

LOADPLUS™ for DB2® General Information

Version 5.1

June 15, 2000



Copyright © 1990–2000 BMC Software, Inc., as an unpublished work. All rights reserved.

BMC Software, the BMC Software logos, and all other BMC Software product or service names are registered trademarks or trademarks of BMC Software, Inc. IBM and DB2 are registered trademarks of International Business Machines Corp. SyncSort is a registered trademark of Syncsort, Inc. All other registered trademarks or trademarks belong to their respective companies.

Restricted Rights Legend

U.S. GOVERNMENT RESTRICTED RIGHTS. UNPUBLISHED–RIGHTS RESERVED UNDER THE COPYRIGHT LAWS OF THE UNITED STATES. Use, duplication, or disclosure by the U.S. Government is subject to restrictions set forth in FAR Section 52.227-14 Alt. III (g)(3), FAR Section 52.227-19, DFARS 252.227-7014 (b), or DFARS 227.7202, as amended from time to time. Send any contract notices to Contractor/Manufacturer:

BMC Software, Inc.
2101 CityWest Blvd.
Houston TX 77042-2827
USA

Contacting BMC Software

You can access the BMC Software Web site at <http://www.bmc.com>. From this Web site, you can obtain general information about the company, its products, special events, and career opportunities. For a complete list of all BMC Software offices and locations, go to <http://www.bmc.com/corporate/offices.html>.

USA and Canada

Address	BMC Software, Inc. 2101 CityWest Blvd. Houston TX 77042-2827
Telephone	713 918 8800 or 800 841 2031
Fax	713 918 8000

Outside USA and Canada

Telephone	(01) 713 918 8800
Fax	(01) 713 918 8000

Customer Support

You can obtain technical support by using Response Online™ (comprehensive information from the Web) or Response On Demand™. To expedite your inquiry, please see “Before Contacting BMC Software,” below.

Response Online

You can obtain technical support from BMC Software 24 hours a day, seven days a week by accessing the technical support Web site at <http://www.bmc.com/support.html>. From this site, you can

- read overviews about support services and programs that BMC Software offers
- find the most current information about BMC Software products
- search a database for problems similar to yours and possible solutions
- order or download product documentation
- report a problem or ask a question
- subscribe to receive e-mail notices when new product versions are released
- find worldwide BMC Software support center locations and contact information, including e-mail addresses, fax numbers, and telephone numbers

Response On Demand

In the USA and Canada, if you need technical support and do not have access to the Web, call 800 537 1813. Outside the USA and Canada, please contact your local support center or your local sales office for assistance.

Before Contacting BMC Software

Before you contact BMC Software, have the following information available so that a technical support analyst can begin working on your problem immediately:

- product information
 - product name
 - product version (release number)
 - license number and password (trial or permanent)
- operating-system and environment information
 - machine type
 - operating system type, version, and service pack or program temporary fix (PTF)
 - system hardware configuration
 - serial numbers
 - related software (database, application, and communication) including type, version, and service pack or PTF
- sequence of events leading to the problem
- commands and options that you used
- messages received (and the time and date that you received them)
 - product error messages
 - messages from the operating system, such as `file system full`
 - messages from related software

Contents

About This Book	xiii
What LOADPLUS Is and Does	1
LOADPLUS Compared to IBM's DB2 LOAD Utility	3
LOADPLUS Features and How They Benefit You	5
LOADPLUS Saves Time and Resources	7
LOAD RESUME PART REPLACE: LOADPLUS V5 and DB2 V6 LOAD	8
LOAD REPLACE: LOADPLUS V5 and DB2 V6 LOAD	9
LOAD REPLACE of Multi-Table Table Space: LOADPLUS V5 and DB2 V6 LOAD	10
Processing Examples	11
Example 1	11
Example 2	12
Example 3	13
Example 4	14
Example 5	15
LOADPLUS Date and Time Formats	17
LOADPLUS Data Type Conversions	19
How LOADPLUS Works	21
The LOADPLUS Operating Environment	22
LOADPLUS Installation	22
How LOADPLUS Achieves Maximum Performance	22
Single-Phase Load	22
Use of DB2 Resources	23
Multiprocessing Environment	23
Use of Virtual Storage	23
Use of Multiple Data Sets	24
Buffering and I/O Optimization	24
Performance Tuning Options	24
The LOADPLUS Execution Phases	25
The LOADPLUS Data Sets	29

Authorization Needed to Execute LOADPLUS	32
Authorization Mechanisms	32
DB2 Authority	32
System (RACF) Authority	33
Single-Phase Load Considerations	34
Referential Integrity and CHECK CONSTRAINTS	35
Recoverability of the Loaded Table Space	36
Hardware Compression Issues	36
 Other BMC Software Products for DB2 and OS/390	37
ACTIVITY MONITOR	38
ALTER	39
Apply Plus	41
APPTUNE	42
CATALOG MANAGER	43
CHANGE MANAGER	44
CHECK PLUS	45
CM/PILOT	46
COPY PLUS and C+/MODIFY	48
Cross-System Image Manager	50
DASD MANAGER PLUS	51
DATA PACKER	52
EXTENDED BUFFER MANAGER	53
OPERTUNE	54
PACLOG	55
RECOVER PLUS	56
R+/CHANGE ACCUM	57
R+/RESOURCE MAXIMIZER	58
RECOVERY MANAGER for DB2	59
RECOVERY MANAGER for OS/390	61
REORG PLUS	63
RESOLVE Log Master for DB2	65
UNLOAD PLUS	66
UTILITY MANAGER	68
 Index	69

Figures

Figure 1	LOAD RESUME PART REPLACE of a Partitioned Table Space	8
Figure 2	LOAD REPLACE of a Partitioned Table Space	9
Figure 3	LOAD RESUME of a Simple Table Space	10
Figure 4	Two-Phase Load Processing Phases	27
Figure 5	Single-Phase Load Processing Phases	28

Tables

Table 1	Comparison of LOADPLUS to IBM's DB2 LOAD Utility	3
Table 2	LOADPLUS Features and Benefits	5
Table 3	Date Formats	17
Table 4	Time Formats	17
Table 5	Timestamp Formats	18
Table 6	Allowable Data Type Conversions	19
Table 7	Minimum Levels of Authority to Access and Update	33

About This Book

This manual provides background and general information about the LOADPLUS for DB2 product from BMC Software. It explains what LOADPLUS does, when it is needed, and how users can benefit from its use.

This document is intended for DB2 system administrators, DB2 database administrators, DB2 application programmers, and corporate officers involved in evaluating prospective software products. It is the best starting point for individuals desiring an overview of the LOADPLUS product, and is one in a complete library of documentation for LOADPLUS and the BMC Software DB2 utility products.

To use this book, you should be familiar with the following items:

- IBM® DB2 Universal Database for OS/390 (DB2)
- Multiple Virtual Storage (MVS) or OS/390 operating systems
- job control language (JCL)

How This Book Is Organized

This book is organized as follows. In addition, an index appears at the end of the book.

Section Number and Title	Description
Section 1: "What LOADPLUS Is and Does"	presents the features of LOADPLUS and explains how they benefit DB2 users. It also compares LOADPLUS to the IBM DB2 LOAD utility and shows benchmark test results. The section provides examples and concludes with tables that indicate the data type conversions supported by LOADPLUS.
Section 2: "How LOADPLUS Works"	provides information on how LOADPLUS works. First the chapter presents the required operating environment and describes how LOADPLUS achieves maximum performance. Next, the section describes and illustrates the LOADPLUS processing phases and data sets. The section concludes with considerations of items such as serialization and concurrency, single-phase load, recoverability, and hardware compression.
Section 3: "Other BMC Software Products for DB2 and OS/390"	describes other products for the DB2 environment available from BMC Software.

Related Documentation

BMC Software products are supported by several types of documentation:

- online and printed books
- online Help
- release notes and other notices

Note: LOADPLUS also provides online message information. For details about how to view the messages that LOADPLUS generates, see the messages section of the *LOADPLUS for DB2 Reference Manual*.

In addition to this book, you can find useful information in the publications listed in the following table. As “Online and Printed Books” on page xv explains, these publications are available on request from BMC Software.

Category	Document	Description
installation	<i>Utilities for DB2 Installation Guide</i>	provides information about using the BMC Software Install System to customize your installation of LOADPLUS and other BMC Software utilities for DB2
	Install System LOADPLUS options panels	provides help for fields required when installing the LOADPLUS product
product-specific documents	<i>LOADPLUS for DB2 Reference Manual</i>	provides detailed information about the LOADPLUS product
	release notes, flashes, technical bulletins	explain the latest updates to LOADPLUS
related documents	<i>Utilities for DB2 Reference Summary</i>	provides a quick reference for commands and syntax for the BMC Software CHECK PLUS for DB2, LOADPLUS, REORG PLUS for DB2, and UNLOAD PLUS® for DB2 products
	<i>BMCDN Command Processor Reference Manual</i>	provides information about the BMC Software utility command processor, BMCDN

Online and Printed Books

The books that accompany BMC Software products are available in online format and printed format. If you are a Windows or Unix user, you can view online books with Acrobat Reader from Adobe Systems. The reader is provided at no cost, as explained in “To Access Online Books.” You can also obtain additional printed books from BMC Software, as explained in “To Request Additional Printed Books.”

To Access Online Books

Online books are formatted as Portable Document Format (PDF) files. You can view them, print them, or copy them to your computer by using Acrobat Reader 3.0 or later. You can access online books from the documentation compact disc (CD) that accompanies your product or from the World Wide Web.

In some cases, installation of Acrobat Reader and downloading the online books is an optional part of the product-installation process. For information about downloading the free reader from the Web, go to the Adobe Systems site at <http://www.adobe.com>.

To view any online book that BMC Software offers, visit the support page of the BMC Software Web site at <http://www.bmc.com/support.html>. Log on and select a product to access the related documentation. (To log on, first time users can request a user name and password either by registering at the support page or by contacting a BMC Software sales representative.)

To Request Additional Printed Books

BMC Software provides printed books with your product order. To request additional books, go to <http://www.bmc.com/support.html>.

Release Notes and Other Notices

Printed release notes accompany each BMC Software product. Release notes provide current information such as

- updates to the installation instructions
- last-minute product information

In addition, BMC Software sometimes provides updated product information between releases (in the form of a flash or a technical bulletin, for example). The latest versions of the release notes and other notices are available on the Web at <http://www.bmc.com/support.html>.

Conventions

This book uses the following general conventions:

Item	Example
information that you are instructed to type	Type SEARCH DB in the designated field. Type search db in the designated field. (Unix)
specific (standard) keyboard key names	Press Enter .
field names, text on a panel	Type the appropriate entry in the Command field.
directories, file names, Web addresses	The BMC Software home page is at www.bmc.com .
nonspecific key names, option names	Use the HELP function key. KEEPDICTIONARY option
MVS calls, commands, control statements, keywords, parameters, reserved words	Use the SEARCH command to find a particular object. The product generates the SQL TABLE statement next.

This book uses the following types of special text:

Note: Notes contain important information that you should consider.

Warning! Warnings alert you to situations that could cause problems, such as loss of data, if you do not follow instructions carefully.

Tip: Tips contain useful information that may improve product performance or that may make procedures easier to follow.

What LOADPLUS Is and Does

DB2 is a powerful relational database management system whose wide success in the DBMS community rests on its ability to provide fast application development and easy data access to the user. As users depend more and more on DB2 for critical business applications, the need for fast operations becomes crucial.

In a DB2 environment, there often is a need to load large amounts of data into DB2 tables. Data is loaded to initially populate a table, periodically replace the data in a table, periodically add data to a table, or reload data into a table whose definition has been changed by a product like the BMC Software CHANGE MANAGER product.

This data can originate from other DB2 tables (possibly from a different DB2 system), from other database management systems (DBMSs) such as IMS, or from applications that generate sequential files. In many cases, the number of rows loaded can amount to millions. And because the data in the table space is unavailable to applications during the loading process, it is important that the loading process be completed as quickly as possible.

The BMC Software LOADPLUS for DB2 product is a reliable, flexible utility that answers the need for accessible and available DB2 data; specifically, it helps to provide fast load operations.

LOADPLUS replaces most of the functions of the IBM DB2 LOAD utility and provides many additional functions that were previously not available. Additionally, LOADPLUS lets users control specific aspects of the load process.

The IBM DB2 LOAD utility loads data from sequential files into tables; however, when you load large quantities of data, the tables cannot be used for a considerable period of time during the LOAD process. The BMC Software LOADPLUS utility shortens the elapsed time of this process.

In addition, LOADPLUS provides solutions for the shortcomings of the IBM LOAD utility in the area of data conversions. Moreover, LOADPLUS eliminates the need to perform additional tasks, such as presorting the input data, making image copies, and gathering DB2 object statistics, before and after you load data.

The BMC Software product line for DB2 responds to the needs of database administrators, system administrators and other DB2 users by providing high-performance database administration and utility products. LOADPLUS is a high-performance utility for loading DB2 data. This product offers a powerful SELECT-like syntax for specifying the data to be loaded, a comprehensive set of data type conversions, and expanded functions.

To summarize, LOADPLUS offers you these benefits:

- **Reduced costs** of loading DB2 data, because fewer CPU cycles and EXCPs are required.
- **Increased availability** of DB2 data, because multiple CPUs, multitasking, and parallel processing reduce the time needed to load the data.
- **Availability of a comprehensive set of data conversions**, which eliminates the need to write special application code.
- **Expands functionality** to eliminate the need for users to perform additional tasks (currently required by DB2) before and after data is loaded

LOADPLUS Compared to IBM's DB2 LOAD Utility

The following table compares the functions performed by the BMC LOADPLUS utility with those available in version 6 of the IBM DB2 LOAD utility.

Table 1 Comparison of LOADPLUS to IBM's DB2 LOAD Utility (Part 1 of 2)

Feature	LOADPLUS	IBM's LOAD UTILITY
Loads data into one or more tables or partitions	*	*
Adds to or replaces existing data	*	*
Allows you to concatenate separate input records	*	*
Accepts input from sequential files	*	*
Accepts FORMAT UNLOAD data	*	*
Allows you to conditionally select data to be loaded	*	*
Allows you to conditionally assign a null or default value to a column	*	*
Provides restart capabilities	*	*
Allows you to update indexes when a small number of new rows are added to an existing table having a large number of rows (INDEX UPDATE)	*	*
Supports the PREFORMAT option	*	*
Provides an option to load ASCII or EBCDIC data	*	*
Offers you the option to not set the COPY pending status on static tables	*	*
Optionally creates inline image copies	*	*
Optionally allocates work files and image copy data sets dynamically	*	sort work files only
Allows you to rebuild indexes, in effect reorganizing the index	*	
Executes concurrently with other BMC Software utilities on DB2 spaces	*	
Allows you to multitask, using multiple SYSREC and SORTOUT data sets	*	
Provides a comprehensive set of data type conversions	*	
Verifies that the data is correct <i>before</i> replacing or adding it to the existing data	*	
Produces BMCSTATS in a statistics report (if the BMC Software DASD MANAGER PLUS product is installed)	*	
Optionally updates DB2 catalog statistics	*	*

Table 1 **Comparison of LOADPLUS to IBM's DB2 LOAD Utility (Part 2 of 2)**

Feature	LOADPLUS	IBM's LOAD UTILITY
Optionally deletes and redefines user-defined (VCAT-defined) data sets as part of the LOAD process	*	
With a self-contained copy function, produces up to four image copies or DSN1COPYs as part of the LOAD process with DB2	*	
Reports multiple errors on the LOAD command you specify	*	
Reports defaulted start and end positions of fields	*	
Provides an option to pause before loading data if any records are discarded	*	
Provides extensive set of comparison operators for WHEN and NULLIF/DEFAULTIF	*	
Allows AND, OR, NOT and () in WHEN conditions	*	
Allows you to use numeric data in WHEN and NULLIF/DEFAULTIF comparisons	*	
Allows you to use the WHEN option to conditionally select FORMAT UNLOAD data	*	
Allows you to load data archived by the BMC Software REORG PLUS product	*	
Allows you to specify the DEFAULTIF default value	*	
Allows you to assign NULL or a default value to a column if a conversion error occurs	*	
Restores PCTFREE and FREEPAGE space in nonleaf pages of indexes	*	
Checks for work file validity and integrity	*	
Maintains a history of LOAD executions	*	
Provides an option for ordering data by table or clustering key if sorting	*	
Optionally allows you to indicate whether all field names should match column names	*	
Analyzes resources for the specific load	*	
Provides data about resource requirements	*	
Optionally allows you to bypass sorting data and the clustering index when the data is already in clustering index order	*	
Offers an optional single processing phase that eliminates work data set requirements and provides additional performance gain in most cases	*	
Allows you to alter the index type while loading the data (ALTERINDEX option)	*	
Optionally sorts input data by clustering key.	*	

LOADPLUS Features and How They Benefit You

The following table details the features of LOADPLUS and the benefits they provide for you.

Table 2 **LOADPLUS Features and Benefits (Part 1 of 2)**

Features	Benefits
Significantly reduces elapsed times, CPU cycles, and EXCPs.	Increases availability of DB2 data. Reduces the cost of loading data.
Takes advantage of multiple CPUs, multitasking, and parallel processing.	Shortens elapsed times for loading data. Increases availability of DB2 data.
Produces BMCSTATS concurrently with the LOAD process and optionally updates statistics in the DB2 catalog.	Eliminates the need for you to execute the RUNSTATS utility separately.
With a self-contained copy function, produces image copies or DSN1COPYs as part of the LOAD process. Produces up to two local and two remote copies.	Eliminates the need for you to execute the COPY utility or DSN1COPY program separately. Fully supports the DB2 recovery process.
Optionally sorts input data by clustering key.	Eliminates the need for you to sort input data separately. Shortens the elapsed time for sorting and loading data.
Provides a comprehensive set of data type conversions.	Eliminates the need for writing special application code to provide data in the allowable input formats. Eliminates the need for you to preprocess data before it is loaded.
Provides more powerful NULLIF and DEFAULTIF processing.	Allows you more control in determining if a column receives NULL or a default value, and which default value it receives.
Creates organized indexes when you use LOAD RESUME YES or LOAD PART REPLACE.	Eliminates the need for you to reorganize your indexes after the data is loaded.
Verifies that the data is correct before replacing or adding it to existing data.	Retains existing data until it is actually replaced by the new data. Increases the availability of your data. Minimizes the need to recover old data if the load is unsuccessful.
Allows you to ignore specified kinds of discards.	Prevents output of useless discard messages and records.
Checks work files for validity and integrity.	Prevents loading of incorrect data due to JCL errors.
Provides installation options to control the use of system resources and sorting options.	Allows you to adapt LOADPLUS for specific needs in your environment.
Optionally provides detailed performance statistics.	Allows you to identify and address specific performance problems.
Maintains a history table of statistics on all previous LOADPLUS activities.	Allows effective planning for future load processing.
Allows you to load data archived by the BMC Software REORG PLUS for DB2 product.	Allows you to process purged information from a REORG PLUS reorganization job.

Table 2 **LOADPLUS Features and Benefits (Part 2 of 2)**

Features	Benefits
Integrates with the BMC Software CHANGE MANAGER product.	Speeds the process of changing or migrating DB2 data structures.
Provides more powerful comparisons and conditions.	Eliminates the need for you to preprocess the data before it is loaded.
Provides full boolean logic.	Allows selection on more than one condition.
Optionally provides ordering by table or clustering key.	Allows similar data from multiple tables in a multitable table space to be placed next to each other for fast access.
Recalls data sets migrated by a storage management system if they are needed for load processing.	Eliminates the need to recall the data set manually before submitting the load job.
Supports partition level independence for objects without secondary indexes.	Allows executing against different partitions of the same table at the same time.
Executes concurrently with other BMC Software utilities on DB2 spaces.	Reduces time required for more than one utility to access different partitions.
Optionally updates indexes rather than rebuilding them with RESUME YES.	Eliminates the time used to unload and sort existing indexes before a load.
Provides multitasking capability by allowing parallel loading of multiple partitions of a partitioned table.	Significantly reduces elapsed times for loading data.
Optionally verifies any DB2 table check constraints existing on the table.	Maintains table check integrity during the load, which avoids having the table space in CHECK pending status.
Estimates resource consumption and provides more information for allocation of work data sets.	Allows more efficient allocation of resources and reduces the possibility of load failure.
Provides the option to dynamically allocate work file and copy data set space.	Eliminates the need to compute work file sizes and determine requirements. Automatically calculates optimal work file sizes and allocates the files.
Provides an option to bypass sorting both the data and the clustering index when the data is already in clustering index order.	Eliminates the time used to sort the data and the clustering index.
Offers single-phase processing.	Eliminates the requirement for work data sets and significantly improves performance in most cases.
Provides a command option to specify the encoding scheme of the input data (ASCII/EBCDIC).	Allows you to load data in either ASCII or EBCDIC format.
Provides the capability to change the index type while loading the data.	Allows you to effortlessly change type 1 indexes to type 2 and vice versa.
Uses the BMC Software sort engine.	Improves sort performance.

LOADPLUS Saves Time and Resources

LOADPLUS completes faster and uses fewer CPU cycles than the IBM LOAD utility. Benchmark tests compared the elapsed time and CPU usage of the following types of loads:

- LOAD RESUME of a partitioned table space, replacing two parts
- LOAD REPLACE of a partitioned table space
- LOAD REPLACE of a multi-table table space in a SAP environment

Results of the comparisons, as illustrated on the following pages, demonstrate that LOADPLUS uses less time and fewer resources to perform the load process.

Note: The benchmark tests using BMC and IBM utilities were run at off hours and on weekends during periods of low CPU activity. The jobs were run several times to ensure repeatability of results. Although BMC Software personnel reviewed the performance measurements to ensure accuracy, *BMC Software does not warrant that similar results will be obtained in any third-party environment.* Your workload and environment may substantially affect your results.


LOAD RESUME PART REPLACE: LOADPLUS V5 and DB2 V6 LOAD

Figure 1 presents time comparisons for loading two parts of a partitioned table space that contain 9.6 million rows of unsorted data.

Version 5 of LOADPLUS shows the following improvements over the DB2 version 6 LOAD utility:

- completed the job nearly **3** times faster
- used **70** percent less CPU time

Figure 1 **LOAD RESUME PART REPLACE of a Partitioned Table Space**




LOAD REPLACE: LOADPLUS V5 and DB2 V6 LOAD

Figure 2 presents time comparisons for loading of a partitioned table space with 28.8 million rows of data. In this scenario, LOADPLUS uses its internal sort engine to sort the data during the load process and the DB2 LOAD utility uses SyncSort to sort the data in a separate step.

Version 5 of LOADPLUS shows the following improvements over the DB2 version 6 LOAD utility:

- completed the job nearly **2** times faster
- used nearly **50** percent less CPU time

Figure 2 **LOAD REPLACE of a Partitioned Table Space**




LOAD REPLACE of Multi-Table Table Space: LOADPLUS V5 and DB2 V6 LOAD

Figure 3 illustrates time comparisons for loading a multi-table table space in a typical SAP environment. This table space contains 97 tables and 145 indexes and the job loads 4,372,520 rows. In this scenario, LOADPLUS uses its internal sort engine to sort the data during the load process and the DB2 LOAD utility uses SyncSort to sort the data in a separate step.

Version 5 of LOADPLUS shows the following improvements over the DB2 version 6 LOAD utility:

- completed the job over **3** times faster
- used over **90** percent less CPU time

Figure 3 **LOAD RESUME of a Simple Table Space**



Processing Examples

The LOADPLUS command syntax provides flexibility to control the loading process to meet your specific needs. Following are five examples that illustrate:

- compatibility between the DB2 LOAD utility and LOADPLUS (Example 1)
- LOADPLUS data conversion and statistics options (Example 2)
- additional data conversions plus image copy capabilities (Example 3)
- use of the PRELOAD PAUSE option (Example 4)
- loading multiple partitions and using the VALUE clause (Example 5)

Example 1

LOADPLUS allows you to use your existing IBM DB2 LOAD utility control statements with little or no modification.

The following example runs with either load utility. Note that LOADPLUS uses the same default data types as the DB2 LOAD utility.

```
LOAD DATA REPLACE
  DISCARDS 25
  INTO TABLE TEST.TBL1
    ( SMINT_COL1      POSITION(1)      INTEGER EXTERNAL(3)
    ,TIMESTAMP_COL2  POSITION(4)
    ,DATE_COL3       POSITION(30)
    ,VARCHAR_COL4    POSITION(40)     CHARACTER(200)
    ,CHAR_COL5       POSITION(240)
    )
  LOG NO
```

Example 2

This example is designed to show some of the data conversion flexibility that LOADPLUS allows.

In addition to allowing you to convert CHAR input data to VARCHAR, LOADPLUS provides an option to TRIM the trailing blanks from the input field.

LOADPLUS permits a wide variety of time and date input formats. This example illustrates this with the following:

- `TIMESTAMP_COL2` is loaded with a timestamp that is already in internal DB2 format on the input file.
- `DATE_COL3` is loaded with date values in YYMMDD format.
- `TIME_COL4` input data is in the format HH.MM.SS.

LOADPLUS allows you to load numeric columns from character fields and designate a default value when data in the input file is not numeric. In this example, the following occurs:

- `SMINT_COL1` is set to NULL when there is a conversion error.
- A record with a non-numeric `SMINT_COL1` input value is not discarded.

In this example, LOADPLUS collects and stores statistics in the DB2 catalog and in the BMCSTATS tables (if the BMC DASD MANAGER PLUS product is installed).

```
LOAD REPLACE
  UPDATEDB2STATS YES BMCSTATS YES
  INTO TABLE TEST.TBL2
    ( SMINT_COL1      POSITION(1)      CHARACTER(3)
      NULLIF ERROR
    ,TIMESTAMP_COL2  POSITION(4)      TIMESTAMP
    ,DATE_COL3       POSITION(14)     DATE-2 EXTERNAL
    ,TIME_COL4       POSITION(20)     TIME-D1 EXTERNAL
    ,VARCHAR_COL4    POSITION(18)     CHARACTER(200) TRIM
    ,CHAR_COL5       POSITION(218)    CHARACTER
    )
```

Example 3

This example shows additional data conversion techniques as well as image copy creation. Because the example specifies RESUME YES, LOADPLUS loads new rows while keeping existing rows intact.

The input file contains two types of records: 'D' (detail) records and 'T' (totals) records. This example loads the detail records.

In this example, the numeric columns, SALES_DOLLAR and SALES_VOL, carry dollars and volumes in 1000s, but the input data carries them in units. The scale of 3 in the field definition causes division by 1000 if the column definition has a scale of 0. The result is rounded up.

LOADPLUS translates the input field for SALES_CODE using NULLIF and DEFAULTIF clauses. These are processed first to last so that if SALES_DATE is less than 01/01/1989, the value put in SALES_CODE is null. If SALES_DATE >= 01/01/1989, then LOADPLUS interrogates the value in the input field SALES_CODE. The column is null if the input field contains an invalid value. A value of '1' translates to 'Y' and '0' translates to 'N'.

LOADPLUS loads SALES_DATE from an input field that contains dates in the format YYYYDDD.

In addition, LOADPLUS makes an image copy and registers it in the DB2 catalog.

```
LOAD RESUME YES
COPY YES REGISTER ALL
INTO TABLE TEST.SALES_DETAIL
  WHEN REC_TYPE = 'D'
    (REC_TYPE          POSITION(1:1)    CHARACTER
    ,PRODUCT_ID        POSITION(2:5)
    ,SALES_DATE         POSITION(6:12)   DATE-3E EXTERNAL
    ,SALES_DOLLAR       POSITION(13:27)  DECIMAL EXTERNAL(15,3) ROUND
    ,SALES_VOL          POSITION(28:36)  INTEGER EXTERNAL(,3)   ROUND
    ,SALES_CODE         POSITION(37:37)
      NULLIF SALES_DATE < '01/01/1989'
      NULLIF SALES_CODE NOT IN ('Y','N','1','0')
      DEFAULTIF SALES_CODE = '1' VALUE('Y')
      DEFAULTIF SALES_CODE = '0' VALUE('N')
    )
```

Example 4

This load is similar to Example 3 except that both detail and totals tables, which are in the same table space, are loaded at the same time. The input file contains two types of records: 'D' (detail) records and 'T' (totals) records. REC_TYPE is a field on the input file but is not a column in either of the tables.

The REPLACE option ensures that LOADPLUS eliminates all existing rows in this load.

This load stops at the end of the PRELOAD phase after LOADPLUS has diagnosed all the errors, built the new rows, and written the rows to an external work file, but before it has loaded the table space and built the indexes. At a later time, you can RESTART the load job and the load will take place. LOADPLUS checks all the rows during the PRELOAD phase to make sure they do not result in duplicate key violations of any UNIQUE indexes on the table.

```
LOAD REPLACE
  PRELOAD PAUSE
  UNIQUECHECK YES
*
  INTO TABLE TEST.SALES_DETAIL
    WHEN REC_TYPE = 'D'
      (REC_TYPE          POSITION(1:1)    CHARACTER
      ,PRODUCT_ID        POSITION(2:5)
      ,SALES_DATE         POSITION(6:12)   DATE-3E EXTERNAL
      ,SALES_DOLLAR        POSITION(13:27)  DECIMAL EXTERNAL(15,3) ROUND
      ,SALES_VOL           POSITION(28:36)  INTEGER EXTERNAL(,3)  ROUND
      ,SALES_CODE          POSITION(37:37)
        NULLIF SALES_DATE < '01/01/1989'
        NULLIF SALES_CODE NOT IN ('Y','N','1','0')
        DEFAULTIF SALES_CODE = '1' VALUE('Y')
        DEFAULTIF SALES_CODE = '0' VALUE('N')
      )
*
  INTO TABLE TEST.SALES_TOTALS
    WHEN REC_TYPE = 'T'
      (REC_TYPE          POSITION(1:1)    CHARACTER
      ,PRODUCT_ID        POSITION(2:5)
      ,TOT_DOLLAR         POSITION(6:20)   DECIMAL EXTERNAL(15,3) ROUND
      ,TOT_VOL            POSITION(21:30)  INTEGER EXTERNAL(,3)  ROUND
      )
```

Example 5

This load is of a table containing sales data for a company that only does business in states touching the Atlantic Ocean.

This example loads rows into three partitions of a table that already contain data. Partitions one and two retain the existing rows but partition three has its existing rows deleted prior to adding new ones.

The high-order part of the clustering key of this table is SALES_REGION. SALES_REGION is not carried on the input file but is built based on the value in SALES_PRSN_STATE from the input file.

This LOADPLUS job creates a DSN1COPY of each of the three partitions.

```
LOAD
  COPY YES REGISTER NONE
  RESUME YES
*
  INTO TABLE TEST.SALESMGR_TBL PART 1
    WHEN SALES_PRSN_STATE IN ( 'CN', 'ME', 'RI', 'NH', 'NY', 'NJ', 'MA' )
      ( SALES_PRSN_LNAME      POSITION(1)   CHARACTER(15)
        , SALES_PRSN_FNAME    POSITION(21:30) CHARACTER
        , SALES_PRSN_MI       POSITION(31:31) CHARACTER
        , SALES_PRSN_STATE    POSITION(32:33) CHARACTER
        , TOTAL_PRSN_SALES    POSITION(34:42) DECIMAL EXTERNAL(9,3) ROUND
        , QTY_PROD1           POSITION(43:52) INTEGER EXTERNAL(,3) ROUND
        , QTY_PROD2           POSITION(53:62) INTEGER EXTERNAL(,3) ROUND
        , SALES_REGION        VALUE( 'NORTHEAST' )
        , DATE_OF_LAST_SALE   POSITION(63:70) DATE-3E EXTERNAL
      )
*
  INTO TABLE TEST.SALESMGR_TBL PART 2
    WHEN SALES_PRSN_STATE IN ( 'NC', 'SC', 'VA', 'GA', 'FL', 'DC', 'MD', 'DE' )
      ( SALES_PRSN_LNAME      POSITION(1)   CHARACTER(15)
        , SALES_PRSN_FNAME    POSITION(21:30) CHARACTER
        , SALES_PRSN_MI       POSITION(31:31) CHARACTER
        , SALES_PRSN_STATE    POSITION(32:33) CHARACTER
        , TOTAL_PRSN_SALES    POSITION(34:42) DECIMAL EXTERNAL(9,3) ROUND
        , TOTAL_QTY_PROD1     POSITION(43:52) INTEGER EXTERNAL(,3) ROUND
        , TOTAL_QTY_PROD2     POSITION(53:62) INTEGER EXTERNAL(,3) ROUND
        , SALES_REGION        VALUE( 'SOUTHEAST' )
        , DATE_OF_LAST_SALE   POSITION(63:70) DATE-3E EXTERNAL
      )
```

(continued on next page)

*

```
INTO TABLE TEST.SALESMGR_TBL PART 3 REPLACE
  WHEN SALES_PRSN_STATE NOT IN
    ( 'NC' , 'SC' , 'VA' , 'GA' , 'FL' , 'DC' , 'MD' , 'DE' ,
      'CN' , 'ME' , 'RI' , 'NH' , 'NY' , 'NJ' , 'MA' )
    ( SALES_PRSN_LNAME      POSITION(1)   CHARACTER(15)
      , SALES_PRSN_FNAME    POSITION(21:30) CHARACTER
      , SALES_PRSN_MI       POSITION(31:31) CHARACTER
      , SALES_PRSN_STATE    POSITION(32:33) CHARACTER
      , TOTAL_PRSN_SALES    POSITION(34:42) DECIMAL EXTERNAL(9,3) ROUND
      , TOTAL_QTY_PROD1     POSITION(43:52) INTEGER EXTERNAL(,3) ROUND
      , TOTAL_QTY_PROD2     POSITION(53:62) INTEGER EXTERNAL(,3) ROUND
      , SALES_REGION        VALUE( 'UNKNOWN' )
      , DATE_OF_LAST_SALE   POSITION(63:70) DATE-3E EXTERNAL
    )
```


LOADPLUS Date and Time Formats

LOADPLUS accepts the following date and time formats:

- the same formats that the DB2 LOAD utility accepts
- additional formats as shown in Table 3, “Date Formats,” Table 4, “Time Formats,” and Table 5, “Timestamp Formats,” on page 18

Table 3 Date Formats

Non-delimited				Delimited	
Format Number			Format	Format No.	Format
Char	Int	Dec		Char	
1	I1	P1	MMDDYY	D1	MM/DD/YY
2	I2	P2	YYMMDD	D2	YY/MM/DD
3	I3	P3	YYDDD	D3	YY/DDD
4	I4	P4	YYWWD	D4	YY/WW/D
5	I5	P5	DDMMYY	D5	DD/MM/YY
6	I6	P6	DDDDD		
1E	I1E	P1E	MMDDYYYY	D1E	MM/DD/YYYY
2E	I2E	P2E	YYYYMMDD	D2E	YYYY/MM/DD
3E	I3E	P3E	YYYYDDD	D3E	YYYY/DDD
4E	I4E	P4E	YYYYWWD	D4E	YYYY/WW/D
5E	I5E	P5E	DDMMYYYY	D5E	DD/MM/YYYY
'f' can be any character. The format number appears in place of the 'f' in the DATE-f EXTERNAL data type.					

Table 4 Time Formats

Non-delimited				Delimited	
Format Number			Format	Format No.	Format
Char	Int	Dec		Char	
1	I1	P1	HHMMSS	D1	HH.MM.SS
2	I2	P2	HHMM	D2	HH.MM
				D3	HH.MM xM
'.' can be any character. xM is either AM or PM. The format number appears in place of the 'f' in the TIME-f EXTERNAL data type.					

Table 5 Timestamp Formats

Non-delimited	
Format No.	Format
1	YYMMDDHHMMSS
2	YYMMDDHHMMSSNNNNNN
3	Store clock 8-byte binary
1E	YYYYMMDDHHMMSS
2E	YYYYMMDDHHMMSSNNNNNN
Formats are CHAR only except 3. The format number appears in place of the 'f' in the TIMESTAMP-f EXTERNAL data type.	

LOADPLUS Data Type Conversions

LOADPLUS provides a comprehensive set of data type conversions. Table 6, “Allowable Data Type Conversions,” illustrates the conversions that LOADPLUS supports. LOADPLUS also supports an option for invoking user-written exit routines that perform special data conversions.

Table 6 Allowable Data Type Conversions

Input	COLUMN											
	SMALLINT	INTEGER	DECIMAL	FLOAT	CHARACTER	VARCHAR	GRAPHIC	VARGRAPHIC	DATE	TIME	TIMESTAMP	ROWID
SMALLINT	D	X	X	X	X	X						
INTEGER	X	D	X	X	X	X			X	X		
INTEGER EXTERNAL	X	X	X	X	X	X			X	X		
DECIMAL	X	X	D	X	X	X			X	X		
DECIMAL EXTERNAL	X	X	X	X	X	X			X	X		
DECIMAL ZONED	X	X	X	X	X	X						
FLOAT	X	X	X	D	X	X						
FLOAT EXTERNAL	X	X	X	X	X	X						
CHAR	X	X	X	X	D	X			X	X	X	
VARCHAR	X	X	X	X	X	D			X	X	X	
GRAPHIC							D	X				
GRAPHIC EXTERNAL							X	X				
VARGRAPHIC							X	D				
DATE					X	X			X			
DATE EXTERNAL									D			
DATE-f EXTERNAL					X	X			X			
TIME					X	X				X		

X = allowable conversion
D = allowable conversion and indicates the default input data type
For a date/time column: the default input data format is EXTERNAL. The default input length (or precision and scale for DECIMAL) is the column length (or precision and scale).
For FLOAT data types: LOADPLUS does not support IEEE Binary Floating Point (BFP) format.
For DATE-f, TIME-f, and TIMESTAMP-f EXTERNAL data types: the format number f describes the representation of the input date/time value. If you don't specify the format number, then ISO, USA, EUR, JIS, or LOCAL is assumed.

Table 6 Allowable Data Type Conversions

Input	COLUMN											
	SMALLINT	INTEGER	DECIMAL	FLOAT	CHARACTER	VARCHAR	GRAPHIC	VARGRAPHIC	DATE	TIME	TIMESTAMP	ROWID
TIME EXTERNAL										D		
TIME-f EXTERNAL					X	X				X		
TIMESTAMP					X	X			X	X	X	
TIMESTAMP EXTERNAL									X	X	D	
TIMESTAMP-f EXTERNAL					X	X			X	X	X	
ROWID												D
EXIT	X	X	X	X	X	X	X	X	X	X	X	
<p>X = allowable conversion</p> <p>D = allowable conversion and indicates the default input data type</p> <p>For a date/time column: the default input data format is EXTERNAL. The default input length (or precision and scale for DECIMAL) is the column length (or precision and scale).</p> <p>For FLOAT data types: LOADPLUS does not support IEEE Binary Floating Point (BFP) format.</p> <p>For DATE-f, TIME-f, and TIMESTAMP-f EXTERNAL data types: the format number f describes the representation of the input date/time value. If you don't specify the format number, then ISO, USA, EUR, JIS, or LOCAL is assumed.</p>												

How LOADPLUS Works

This section provides the following information about how the LOADPLUS for DB2 product works:

- the required operating environment and how LOADPLUS is installed
- how LOADPLUS works with DB2 table data while operating outside of the DB2 subsystem
- a description and diagram of the LOADPLUS processing phases and a description of the LOADPLUS data sets
- an explanation of the authorization needed to execute LOADPLUS
- discussions of single-phase load, referential integrity, recoverability, and hardware compression

The LOADPLUS Operating Environment

LOADPLUS runs with DB2 versions 3 and later. It does not support DB2 version 2.3 or earlier versions.

LOADPLUS executes as a batch job and uses the IBM DB2 Call Attach Facility. LOADPLUS also requires the IBM DFSORT program or a compatible system sort routine.

LOADPLUS Installation

You install LOADPLUS by using the BMC Software Install System (DCI). The installation process does not require any modifications to DB2, nor does it require an IPL of your system.

How LOADPLUS Achieves Maximum Performance

LOADPLUS uses advanced processing techniques to provide significant performance advantages over IBM's standard load processing method. These processing techniques allow LOADPLUS to be optimized for a single purpose—to achieve maximum utility performance.

Single-Phase Load

LOADPLUS provides two architectures: two-phase load and single-phase load. When you specify PRELOAD CONTINUE (the default), PRELOAD PAUSE, or PRELOAD ANALYZE, LOADPLUS uses two phases, PRELOAD and LOAD. When you specify PRELOAD LOAD, LOADPLUS combines the PRELOAD and LOAD phases into a single phase called COMBINED.

The two-phase architecture allows LOADPLUS to perform several processes concurrently, reducing the elapsed time for load processing. For example, LOADPLUS checks for duplicate keys while it reads input data and sorts index keys.

The single-phase architecture builds on the advantages of the two-phase processing and allows for even greater reductions in CPU usage and elapsed time under certain conditions due to the reduction of read and write operations. It also eliminates the need to provide work data sets because LOADPLUS writes the data directly to the table space and index.

Use of DB2 Resources

LOADPLUS performs most of its work outside the DB2 subsystem. Therefore, it is not dependent on DB2's general-purpose facilities and services, and its impact on those resources and other types of DB2 work is minimized.

LOADPLUS interacts with DB2 in the following ways:

- uses SQL to access the DB2 catalog
- uses SQL to update the LOADPLUS tables (BMCUTIL, BMCHIST, BMCSYNC, and BMCDICT)
- uses DB2 commands to STOP/START spaces

Multiprocessing Environment

LOADPLUS exploits the technology provided by large-scale processors. During execution, LOADPLUS examines available resources and uses as much of these resources as possible to maximize performance and reduce elapsed time.

During execution, LOADPLUS also determines the most effective arrangement of tasks when running in a multiprocessor environment. Although LOADPLUS executes very efficiently on single-processor computers, it performs even better on large multiprocessor systems.

Use of Virtual Storage

Because each task requires virtual storage for processing, LOADPLUS balances the multiprocessing performed with the amount of virtual storage it finds available. The utility is designed to use as much virtual storage as it needs, providing it is available. If the maximum amount it needs is not available, LOADPLUS uses less.

Use of Multiple Data Sets

LOADPLUS provides the capability to load data from multiple input data sets (SYSRECs) and to use multiple SORTOUT data sets for partitioned tables in the PRELOAD and LOAD phases. This multitasking results in reduced elapsed times for loading your data.

LOADPLUS optimizes the number of sort processes based on several factors, such as the number of SORTOUT data sets you specify, how many sort processes will fit in the available memory, and the number of partitions involved in the process. Many of these factors can be adjusted either in the installation options, the DD statements, or the LOADPLUS command options.

Buffering and I/O Optimization

The single most important factor affecting the performance of LOADPLUS is I/O processing. This is because during a typical execution, LOADPLUS reads and writes large amounts of data.

To maximize I/O performance, LOADPLUS handles all its own buffering and performs I/O operations at the lowest level possible. This allows reading and writing of several blocks of data with each I/O operation and allows LOADPLUS to prefetch subsequent data. LOADPLUS calculates the optimum block size for data sets and uses the largest block size possible to reduce the number of I/O operations.

Performance Tuning Options

LOADPLUS provides global processing options to control such things as the number of buffers used by input and output data sets, parameters to be passed to the sort routine when it is invoked by LOADPLUS, and the maximum level of multitasking allowed. You can change these parameters when necessary to tune LOADPLUS to achieve maximum performance in your operating environment.

The LOADPLUS Execution Phases

The following are the execution phases of LOADPLUS and their primary functions. The diagrams starting on page 27 illustrate what happens during each phase of execution.

UTILINIT	Initializes the job, reads, parses, and verifies the LOAD command and the commands in the SYSIDCIN data set. It also performs DB2 catalog lookup.
ANALYZE	This optional phase analyzes the objects being loaded and optionally produces statistics to help with data set allocations. LOADPLUS executes this phase when dynamic work file allocation is active or when you specify the ANALYZE and ENUMROWS options.
PRELOAD	<p>During the PRELOAD phase, LOADPLUS reads data from the input data sets. It converts and verifies that the data is correct, builds DB2 rows for loading in the SORTOUT data sets, builds index information in the SYSUT1 data sets, and writes any input records in error to the SYSDISC discard data set. It also builds the compression dictionary and compresses the data rows.</p> <p>If you are loading more than one table in a segmented table space, sorting the rows using the ORDER YES option, or if you have unique indexes and are checking for duplicate keys, the PRELOAD phase invokes the sort routine.</p> <p>If you are loading rows into tables or partitions that already contain data, the PRELOAD phase also reads the indexes and table space of the table or partition being loaded.</p>
LOAD	This phase redefines the VSAM data sets when you specify REDEFINE YES. LOADPLUS reads the rows from the SORTOUT and SYSUT1 data sets and loads the data into the table space and indexes. It invokes the sort routine for indexes not sorted in the PRELOAD phase. It creates any requested copies and registers them with the DB2 catalog.
COMBINED	This optional phase combines all the functions of the PRELOAD and LOAD phases into a single processing phase. The COMBINED phase avoids the intermediate steps of writing to the SORTOUT and SYSUT1 data sets by writing data directly to the table space and index space. You do not have to specify these data sets in the JCL, eliminating the need to allocate disk and/or tape resources in large environments.
UTILTERM	Performs cleanup for the utility. This phase sets the CHECK pending status of all affected tables and table spaces, and updates the BMC Software BMCHIST table, the BMCSTATS (if the BMC Software DASD MANAGER PLUS product is installed), and updates the DB2 catalog statistics.

| **All Phases**

All phases update the BMCUTIL and BMCSYNC DB2 tables.

Note: The BMCHIST, BMCUTIL, BMCSYNC, and BMCDICT tables are delivered and installed as part of the LOADPLUS product.

Figure 4 Two-Phase Load Processing Phases

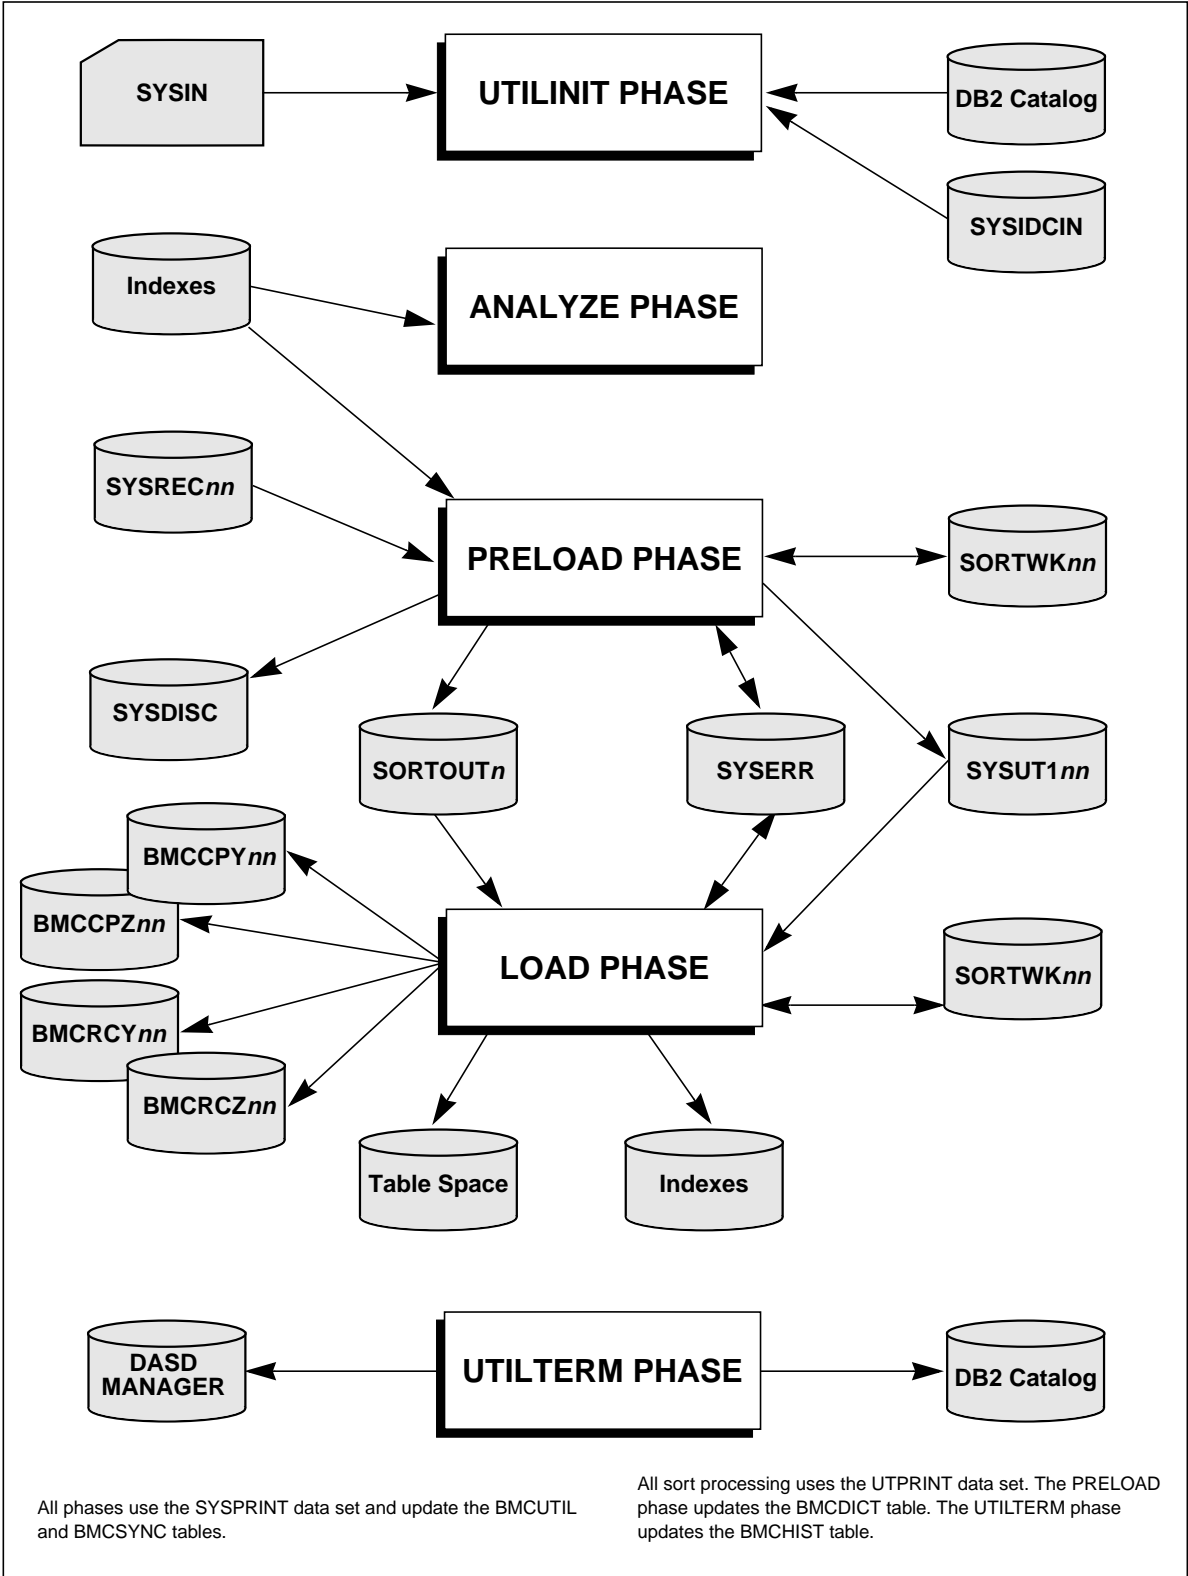
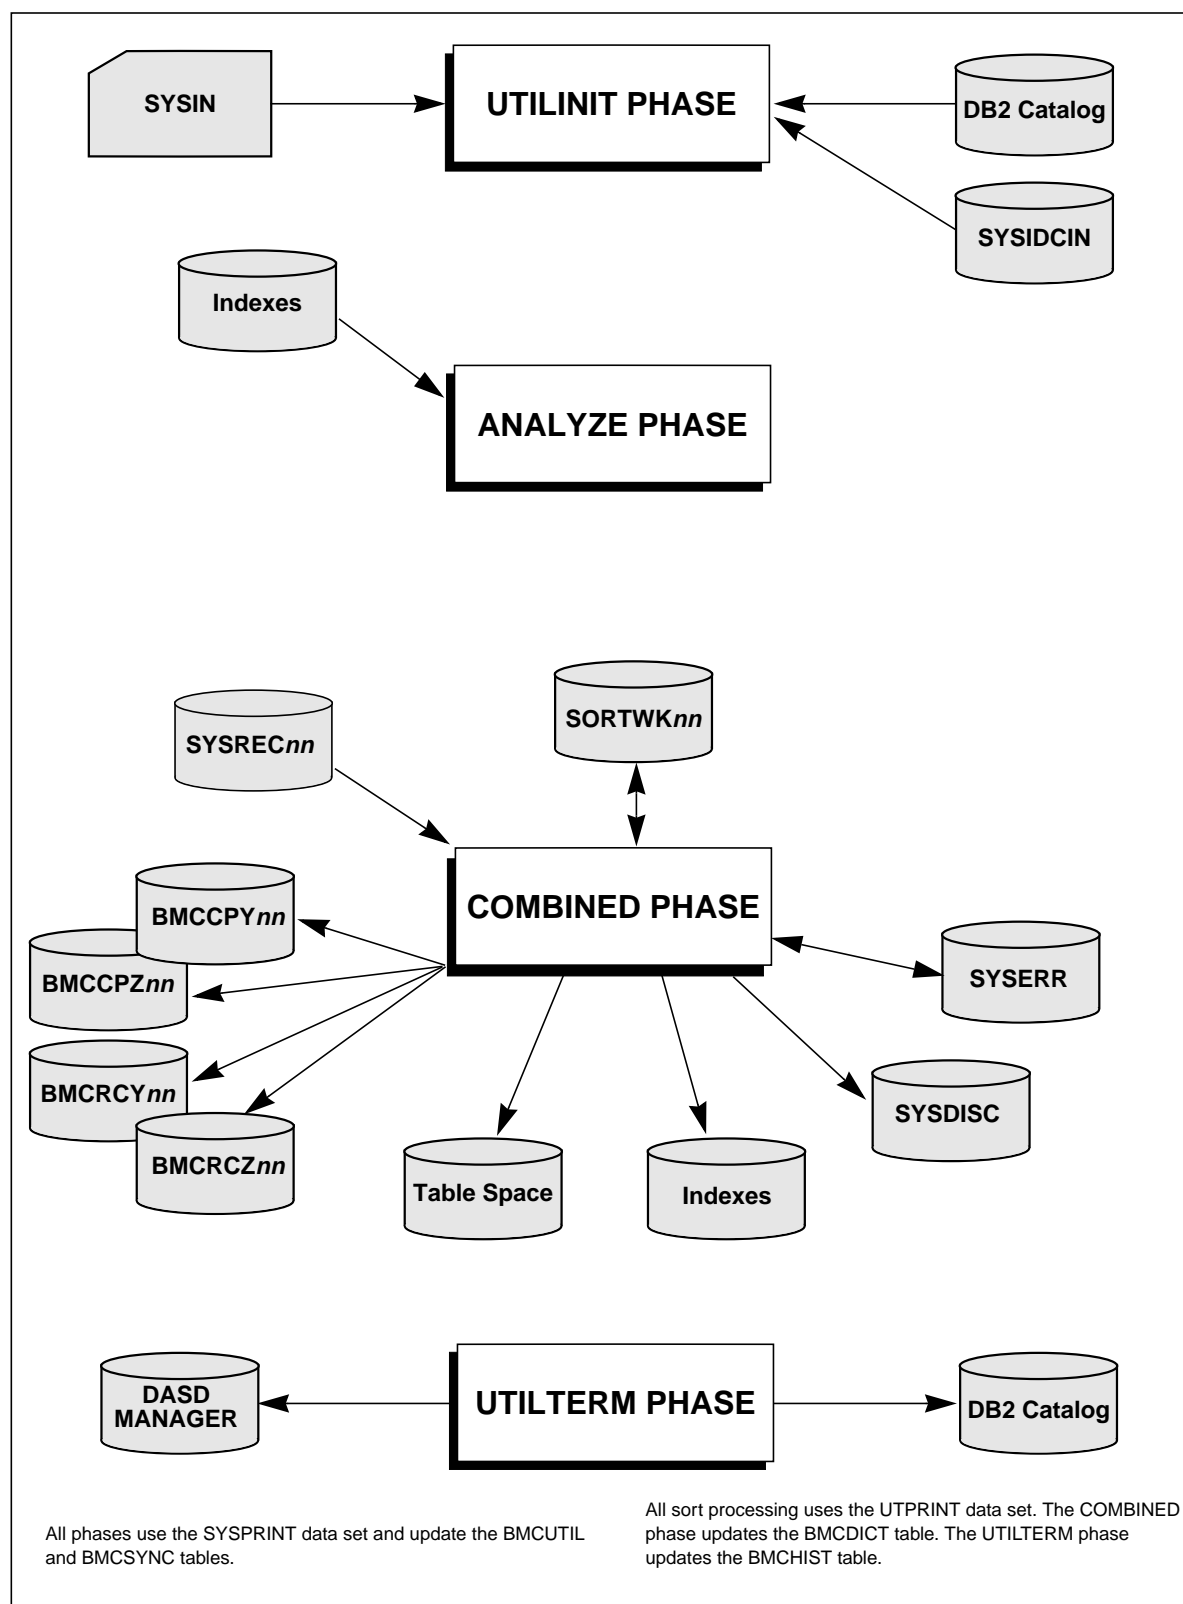


Figure 5 Single-Phase Load Processing Phases



The LOADPLUS Data Sets

LOADPLUS uses the following data sets specified by the ddnames.

SYSIN	The input data set that contains the LOAD command.
SYSRECnn	The input data set containing the data to be loaded. The nn is required only if you specify multiple input data sets. LOADPLUS uses the data set in the PRELOAD phase or COMBINED phase. At the end of the PRELOAD or COMBINED phase, the data set is freed.
SORTWKnn	<p>The work data set used by the sort routine. LOADPLUS uses the data set in the PRELOAD and in the LOAD phases for two-phase load and in the COMBINED phase for single-phase load. The sort work files cannot be allocated as VIO data sets or tape data sets. Each individual work file must be allocated on a single DASD unit and cannot be allocated as multivolume.</p> <p>As an alternative to specifying the SORTWKnn data set in a DD statement, you can have LOADPLUS dynamically allocate it.</p>
SYSUT1nn	<p>For two-phase load, the data set that contains index key entries and becomes the input for the index-building process. The nn is required only if you specify multiple data sets. LOADPLUS uses the data set in the PRELOAD and LOAD phases. The data set is deallocated after the index is built in order to free the device.</p> <p>As an alternative to specifying the SYSUT1nn data set in a DD statement, you can have LOADPLUS dynamically allocate it.</p> <p>For a single-phase load, the data set is not required and is not used if specified.</p>
SORTOUTn	<p>For two-phase load, the data set that contains the output of the PRELOAD phase, and is used as input to the LOAD phase. The n is required only if you specify multiple data sets. LOADPLUS uses the data set during the PRELOAD phase through the end of the LOAD phase. After the table space is loaded, the data set is deallocated in order to free the device.</p> <p>As an alternative to specifying the SORTOUTn data set in a DD statement, you can have LOADPLUS dynamically allocate it.</p> <p>For single-phase load, the data set is not required and is not used if specified.</p>

SYSIDCIN	<p>The input data set that contains the IDCAMS command statements that LOADPLUS uses to redefine user-defined (VCAT-defined) data sets. It is read, parsed, and partially verified in the UTILINIT phase. The IDCAMS commands are issued prior to loading the data set and are used for redefinition of VCAT-defined data sets only.</p>
SYSDISC	<p>The data set that contains the discarded input records. LOADPLUS uses this data set only in the PRELOAD phase for two-phase load or the COMBINED phase for single-phase load. At the end of the PRELOAD or COMBINED phase, the data set is freed.</p> <p>As an alternative to specifying the SYSDISC data set in a DD statement, you can have LOADPLUS dynamically allocate it.</p>
SYSERR	<p>The data set that contains information about errors. LOADPLUS uses this data set in the PRELOAD phase through the end of the LOAD phase for two-phase load or the COMBINED phase for single-phase load.</p> <p>As an alternative to specifying the SYSERR data set in a DD statement, you can have LOADPLUS dynamically allocate it.</p>
BMCCPYnn	<p>The output data set that contains the DSN1COPY or image copy. Depending on the option you choose, LOADPLUS creates either a standard image copy or DSN1COPY <i>after</i> it loads the data into the table, or an inline image copy <i>as</i> it loads the data into the table. After the copy is complete, the data set is deallocated in order to free the device unless VOL=(,RETAIN) is specified. (BMCCPY should not be a temporary data set.)</p> <p>As an alternative to specifying your BMCCPYnn data set in a DD statement, you can have LOADPLUS dynamically allocate the data set.</p>
BMCCPZnn	<p>The output data set that contains an additional copy file. Depending on the option you choose, LOADPLUS creates either a standard image copy or DSN1COPY <i>after</i> it loads the data into the table, or an inline image copy <i>as</i> it loads the data into the table. After the copy is complete, the data set is deallocated in order to free the device unless VOL=(,RETAIN) is specified. (BMCCPZ should not be a temporary data set.)</p> <p>As an alternative to specifying your BMCCPZnn data set in a DD statement, you can have LOADPLUS dynamically allocate the data set.</p>

BMCRCYnn	<p>The output data set that contains the DSN1COPY or image copy. Depending on the option you choose, LOADPLUS creates either a standard image copy or DSN1COPY <i>after</i> it loads the data into the table, or an inline image copy <i>as</i> it loads the data into the table. After the copy is complete, the data set is deallocated in order to free the device unless VOL=(,RETAIN) is specified. (BMCRCY should not be a temporary data set.)</p> <p>As an alternative to specifying your BMCRCYnn data set in a DD statement, you can have LOADPLUS dynamically allocate the data set.</p>
BMCRCZnn	<p>The output data set that contains an additional copy file. Depending on the option you choose, LOADPLUS creates either a standard image copy or DSN1COPY <i>after</i> it loads the data into the table, or an inline image copy <i>as</i> it loads the data into the table. After the copy is complete, the data set is deallocated in order to free the device unless VOL=(,RETAIN) is specified. (BMCRCZ should not be a temporary data set.)</p> <p>As an alternative to specifying your BMCRCZnn data set in a DD statement, you can have LOADPLUS dynamically allocate the data set.</p>
SYSPRINT	<p>The output data set that will contain LOADPLUS messages.</p>
UTPRINT	<p>A data set that indicates that sort messages should be reported. However, the actual messages for each sort appear in separate SYSnnnnn data sets, where nnnnn is a system-assigned sequential number.</p>
DB2 Data Sets	<p>LOADPLUS dynamically allocates the DB2 data sets (table spaces or index spaces) you are loading. Therefore, you do not need to specify them in your JCL.</p>

Authorization Needed to Execute LOADPLUS

Using LOADPLUS requires that you have authorization within DB2 and through your system security package (such as RACF) sufficient to access resources and perform the tasks accomplished during LOADPLUS processing.

Authorization Mechanisms

If you have a DB2 security exit in place (DB2 version 5 or later) as your authorization mechanism, LOADPLUS uses that exit to verify utility authorizations. IBM supplies a sample DB2 security exit, which uses RACF.

If you use the CA-ACF2 for DB2 or CA-Top Secret for DB2 security product from Computer Associates as your authorization mechanism, LOADPLUS uses that product to verify utility authorizations. LOADPLUS can use that product with any version of DB2 after it detects the presence of the product in the subsystem where LOADPLUS is running.

To use CA-ACF2 for DB2 or CA-Top Secret for DB2, you must have applied APAR TA4715D from Computer Associates, and you must specify ACFORTSS=YES in your installation options module.

If it detects none of these mechanisms, LOADPLUS uses the standard DB2 method to check security.

DB2 Authority

To run LOADPLUS, you must have the following authorizations:

- sufficient DB2 authority to execute all LOADPLUS plans
- one of the following authorizations with your primary or secondary authorization ID:
 - SYSADM or SYSCTRL authority
 - DBADM or DBCTRL authority and LOAD authority for the database containing the named tables and associated indexes
 - LOAD authority for the database containing the named tables and their indexes

- LOAD authority for the named tables, or LOAD authority for all the tables of the table space if you specify LOAD REPLACE (LOAD authority means the table's creator ID matches your authorization ID or an ID in your secondary authorization ID list.)
- LOAD authority and ALTER INDEX authority for the database containing the named table space or index

System (RACF) Authority

Because LOADPLUS does not run as part of the DB2 subsystem, you must have system authorization equivalent to the authorization required by DB2 in order to use LOADPLUS. If RACF (or a similar system security package) protects both the underlying data sets and the ICF (Integrated Catalog Facility) catalog of a table or index space, you must have the minimum levels of authority as shown in Table 7.

Note: If LOADPLUS was installed with option OPNDB2ID=YES, these RACF authorities are not required for the user running LOADPLUS.

Table 7 Minimum Levels of Authority to Access and Update

Table or Index Space Definition	To Access and Update DB2 Data Sets	To Access and Update the ICF Catalog
VCAT-defined	CONTROL	UPDATE
STOGROUP-defined	ALTER or CONTROL	UPDATE or CONTROL

Single-Phase Load Considerations

You can choose to combine the PRELOAD and LOAD phases of LOADPLUS into a single processing phase, named COMBINED, by specifying PRELOAD LOAD. The COMBINED phase provides all the capabilities, features, and functions of the PRELOAD and LOAD phases for the following load types:

- LOAD REPLACE
- LOAD RESUME NO
- LOAD RESUME YES PART REPLACE, which you can only use if you are replacing all of the partitions involved in the load and there are no nonclustering indexes

When you specify PRELOAD LOAD, LOADPLUS does not require the SORTOUT and SYSUT1 data sets, and does not use them if they are specified. This provides a significant resource savings when processing large amounts of data, and eliminates all EXCP processing associated with these data sets.

In most cases, single-phase load can provide CPU and elapsed time savings, especially in environments requiring the use of tape for the SORTOUT and SYSUT1 data sets. Tape mount time becomes nonexistent and CPU consumption can be reduced because there is no read/write activity to these data sets.

Note: Two-phase load can be faster if you expect to discard a significant number of duplicate keys. The COMBINED phase loads the data into the table before checking for duplicate keys in the unique non-clustering indexes. If single-phase load detects any unique key violations, it deletes the rows causing the duplicates and the duplicate key entries for *all* indexes for those rows.

The fastest load is a single-phase load with LOAD REPLACE ORDER PRESORTED and no nonclustering indexes. You can specify ORDER PRESORTED on a single-phase load with RESUME YES PART REPLACE if you do all of the following:

- replace all partitions involved in the load
- have no nonclustering indexes
- use only one SYSREC data set

You also should keep in mind that the COMBINED phase of LOADPLUS is a combination of the PRELOAD and LOAD phases. Because the table space being loaded is stopped at the beginning of the phase, the amount of time the objects are unavailable can be longer with single-phase load.

Referential Integrity and CHECK CONSTRAINTS

There is no ENFORCE CONSTRAINTS option for referential integrity violations in LOADPLUS. Instead, LOADPLUS places the table spaces in CHECK pending status if the CHEKPEND=YES option is specified in the installation options member and a table has referential constraints defined. After the LOADPLUS utility completes, you must run CHECK PLUS (specifying the CHECK DATA SCOPE ALL option) or the IBM CHECK DATA utility (specifying SCOPE ALL) to ensure that referential integrity has not been violated.

LOADPLUS allows an image copy to be made after a table is loaded, even when it leaves the table in CHECK pending status. Because either check utility logs any changes made to the table, a full RECOVER utility execution using the copy applies the log records, thus ensuring the integrity of the data in the table. A RECOVER TOCOPY utility execution places the affected table in CHECK pending status again. A subsequent invocation of either check utility restores referential integrity.

If DB2 table check constraints exist for a table being loaded, LOADPLUS verifies all constraints and discard any violations if ENFORCE CHECK CONSTRAINTS is defaulted to or specified. If table check constraints exist, you specify ENFORCE NO, and the CHEKPEND=YES installation option is in effect, LOADPLUS places the table space in CHECK pending status.

After the LOADPLUS utility completes, you must run CHECK PLUS (specifying the CHECK DATA SCOPE ALL option) or the IBM CHECK DATA utility (specifying SCOPE ALL) to ensure that table check integrity has not been violated.

Recoverability of the Loaded Table Space

If you do not take an image copy as part of the LOADPLUS job, you must take some action to ensure that DB2 can recover the table space. If you want an image copy, run the BMC Software COPY PLUS utility or the IBM DB2 COPY utility. You can also execute DSN1COPY or any other suitable substitute your installation provides.

If you intend to create the copy outside of LOADPLUS, your installation options should specify LOADCPY=YES so that LOADPLUS inserts a LOAD LOG(NO) row into SYSIBM.SYSCOPY. This row lets DB2 know that a load has been performed on the table space without logging. Using this row is necessary for recovery to work during fallback processing and to prevent you from making an incremental image copy after a LOAD LOG(NO) but before you make a full image copy.

If you specify COPY NO and your installation options specify LOADCPY=NO, LOADPLUS does not insert a LOAD LOG(NO) row into SYSIBM.SYSCOPY. This action can render the table space unrecoverable or only partially recoverable if the full image copy you make after the load is not usable or if you make incremental image copies after the load but before you make any full image copy.

Hardware Compression Issues

DB2's compression dictionary built by the REORG or LOAD utility and BMC Software's compression dictionary built by LOADPLUS can be used interchangeably.

Like DB2's LOAD, LOADPLUS provides a KEEPDICIONARY option to prevent the utility from building a new compression dictionary. This option is valid only if the table space being loaded has the COMPRESS YES attribute. If a dictionary already exists, LOADPLUS uses it for compression. If a dictionary does not exist, LOADPLUS builds one during the PRELOAD phase for two-phase load and during the COMBINED phase for single-phase load. After it completely builds the dictionary, LOADPLUS compresses the rest of the data. LOADPLUS stores the dictionary in the BMC_DICT table between phases for two-phase load.

The reported percentage of compressed rows does not include rows that are compressed by an EDITPROC.

LOADPLUS reserves enough pages for the largest dictionary size.

Other BMC Software Products for DB2 and OS/390

BMC Software offers a comprehensive line of products for DB2 and OS/390 for use by system administrators, database administrators (DBAs), and developers. These products include a set of integrated database administration tools and fast utilities that provide enhanced functions while decreasing the demand on system resources. In addition to LOADPLUS™, the BMC Software products for DB2 and OS/390 include

- ACTIVITY MONITOR (page 38)
- ALTER® (page 39)
- Apply Plus (page 41)
- APPTUNE™ (page 42)
- CATALOG MANAGER (page 43)
- CHANGE MANAGER (page 44)
- CHECK PLUS (page 48)
- CM/PILOT® (page 46)
- COPY PLUS and C+/MODIFY™ (page 48)
- Cross-System Image Manager (XIM™) (page 50)
- DASD MANAGER PLUS (page 51)
- DATA PACKER® (page 52)
- EXTENDED BUFFER MANAGER (XBM™) (page 53)
- OPERTUNE® (page 54)
- PACLOG™ (page 55)
- RECOVER PLUS (page 56)
- R+/CHANGE ACCUM® (page 57)
- R+/RESOURCE MAXIMIZER® (page 58)
- RECOVERY MANAGER for DB2 (RMGR™) (page 59)
- RECOVERY MANAGER for OS/390 (page 61)
- REORG PLUS (page 63)
- RESOLVE Log Master™ (page 65)
- UNLOAD PLUS® (page 66)
- UTILITY MANAGER (page 68)

ACTIVITY MONITOR

The ACTIVITY MONITOR for DB2 product enables you to monitor the performance of DB2 transactions as they occur in real time. Using the time sharing option (TSO), the Interactive System Productivity Facility (ISPF), or the virtual telecommunications access method (VTAM) for presentation, ACTIVITY MONITOR helps you to

- analyze DB2 data from and issue commands to all active DB2 subsystems in the sysplex
- analyze DB2, IMS, CICS, and MVS data from control blocks and traces
- start, stop, and schedule DB2 traces immediately or in advance
- store trace records in data sets according to the type of records collected
- view historical data online or print historical reports
- identify values that fall outside user-defined limits (exceptions) as such values occur and initiate actions based on exception detection:
 - notify users and issue write-to-operator messages (WTOs)
 - start a trace
 - log reports
 - issue a command
- issue ACTIVITY MONITOR, DB2, MVS, OPERTUNE, and XBM commands and receive responses without leaving ACTIVITY MONITOR
- dynamically EXPLAIN Structured Query Language (SQL) statements as transactions process
- move quickly from summary data to detailed information about the DB2 event that you choose
- customize reports to your own specifications:
 - include the DB2 and ACTIVITY MONITOR data of your choice
 - compute formulas that use full Boolean and conditional logic
 - use a wide (132 columns) or narrow (80 columns) format
 - display data as tables or bar charts

ACTIVITY MONITOR also provides a variety of utility functions to assist the administrator in controlling access to the product and defining the resources that ACTIVITY MONITOR can use.

ALTER

The ALTER for DB2 product from BMC Software provides a powerful solution to the problems of managing your DB2 environment. By automating and simplifying the change process, ALTER enables you to deal effectively with the demands of a constantly changing DB2 environment. For example, ALTER enables you to automate the following common tasks:

- create, modify, and drop data structures within a DB2 subsystem
- create new data structures by using existing data structures as templates
- determine the dependencies on changed objects and preserve those dependencies and their associated data
- migrate data structures and data from one DB2 subsystem to another, or migrate them within a subsystem
- analyze the impact of changes by validating the changes against the DB2 catalog

If ALTER detects conflicts, it notifies you through error messages and warnings.

- implement a least-cost strategy for performing changes

You can group your changes to minimize the work that is required to run them.

- import data definition language (DDL) files from another subsystem and apply the changes to your subsystem
- perform space estimation for table spaces and indexes
- provide passive management of authorizations, plans, and packages
- preserve data across structure changes and migrate data from one DB2 subsystem to another when structures are migrated
- provide automatic data conversion for changes to column attributes
- handle requirements for data set allocations

ALTER provides full management support for the following DB2 data structures:

- storage group
- database
- table space
- table
- view
- check constraint
- foreign key
- index
- synonym
- alias

When you use ALTER to specify changes for any of these data structures, ALTER automatically propagates the changes to any dependent objects. For example, if you change the name of a table, ALTER creates a corresponding change in the indexes, synonyms, and other dependent objects that reference the table under its former name.

ALTER uses BMC Software or DB2 utilities in worklists and elsewhere when needed. By using installed BMC Software utilities instead of DB2 utilities, you can significantly enhance the performance of ALTER. BMC Software utilities run faster, provide additional features, and may reduce the number of steps in a worklist.

You can use the following BMC Software utilities with ALTER:

- COPY PLUS
- UNLOAD PLUS
- LOADPLUS
- DASD MANAGER PLUS
- CHECK PLUS
- RECOVER PLUS
- REORG PLUS
- R+/RESOURCE MAXIMIZER

ALTER is a subset of (and can be easily upgraded to) the CHANGE MANAGER product (discussed on page 44).

Apply Plus

The BMC Software Apply Plus product complements the existing backup and recovery product group. It allows you to rapidly apply updates to an IBM DB2 database by using flexible processing configurations that you define.

Using Apply Plus, you can

- migrate data between active DB2 tables
- perform high-speed transaction recovery (undo or redo SQL processing)
- extend the capabilities of the PATROL DB-Log Master product for data migration and transaction recovery
- enhance your existing batch SQL applications by replacing the utilities that you currently use to apply changes to the target tables

Apply Plus also includes the following management features:

- conflict detection and processing
 - exception handling for timeout conditions, relational integrity errors, and multi-row updates
 - retry capability at both the statement and transaction level
 - the ability to defer statements for later correction and reprocessing
- object name mapping for table and column names
- the ability to restart processing in the event of a system or database failure

Apply Plus can process data from logical log files created by PATROL DB-Log Master, or from SQL files (including SQL generated by PATROL DB-Log Master). To achieve its high processing speed, Apply Plus takes advantage of the following performance benefits:

- multiple parallel connections to the DB2 subsystem that can serve a number of apply agents
- static SQL processing, using dynamic DB2 bind processing
- workload balancing by table and partition, with respect for referential integrity constraints and table groups

APPTUNE

The APPTUNE for DB2 product allows you to capture and display data that helps you tune SQL statements quickly, effectively, and at minimal cost. Using APPTUNE, you can

- capture data for each SQL statement (OPEN, FETCH, CLOSE, and so on) that is run during a collection period (both static and dynamic), and view both as a single work load
- obtain statistics and data at the SQL statement level about accounting, buffer, I/O, SQL text, DB2 object data, and lock activities
- analyze actual DB2 object accesses, including getpages and I/O
- identify and capture detailed data about exceptions, including host variable values
- start, stop, and reset data collection activities while the Data Collector remains active
- measure only the resources that were consumed during the processing of an SQL statement
- dynamically EXPLAIN SQL statements, providing both statistical and textual information about the access path, with suggestions about how to improve SQL statement performance
- identify SQL errors and view the corresponding SQL communication area (SQLCA) data and SQL text
- display current information about a DB2 subsystem and active threads
- specify the time period, DB2 subsystems, and data source (archived data or online data) used to generate online or batch reports
- use your own criteria to define the plans and programs that make up an application group for reporting
- choose whether reports are displayed in a traditional, numeric format or a graphical format
- analyze all of the DB2 subsystems in the sysplex at the same time
- issue APPTUNE, DB2, MVS, OPERTUNE, and XBM commands and receive responses without leaving APPTUNE

APPTUNE also provides a variety of utility functions to assist the administrator in controlling access to the product and defining the resources that APPTUNE can use.

CATALOG MANAGER

The CATALOG MANAGER product greatly facilitates the administration of DB2 databases. Through an ISPF-based interface, CATALOG MANAGER provides interactive access to DB2 catalog information by using simple-to-use menus, dialogs, and online help.

CATALOG MANAGER supports the following actions and functions:

- finds DB2 catalog information by listing DB2 objects, related objects, or objects with certain attributes, and then initiates an action on one or more objects in the list
- connects to any DB2 subsystem by using Distributed Data Facility (DDF)
- creates new DB2 objects by using existing objects as a template with hierarchical data definition language (DDL)
- drops catalog objects with the option of viewing a dependency list and logs dropped object structures and copy information for recovery
- grants, revokes, and copies authorizations for users or objects; lists revoke dependencies and reassigns authorizations
- generates and runs up to eight IBM and BMC Software utilities against one or more objects
- tests host variables by using extended SQL processing
- builds and issues DB2 commands
- automatically logs user actions to provide an audit trail
- converts database request modules (DBRMs) to packages or packages to DBRMs
- uses the real DB2 catalog or a shadow catalog
- provides extensive reports on catalog objects, including DESCRIBE for application development and EXPLAIN for SQL tuning

You do not need a thorough knowledge of DB2 structures or SQL syntax. CATALOG MANAGER constructs the necessary SQL statements from your input and selections and also maintains the database structures. You choose when and how to run the SQL statements. In addition, you can store frequently used SQL statements for later use.

CHANGE MANAGER

CHANGE MANAGER enables the DBA to deal effectively with the demands of a constantly changing environment that involves multiple DB2 subsystems. CHANGE MANAGER includes all of the features of the ALTER product (discussed on page 39) and provides the following additional capabilities:

- migrates data structure changes across multiple databases and subsystems
- determines changes to data structures and migrates those changes to one or more copies of the data structures
- captures and records structure definitions and data within a DB2 subsystem (to establish a baseline)
- recovers structures and data to a point in time defined by a baseline within a DB2 subsystem
- compares two versions of structure definitions to
 - determine the changes necessary to upgrade one version to another
 - selectively apply changes to copies of the data structures while preserving the uniqueness of each copy
- uses CASE tool output to determine the changes to existing DB2 application structures
- reduces the volume of information that is needed to communicate changes by using a BMC Software language called Change Definition Language™ (CDL®) to transmit the change information
- feeds changes that are made on a remote system back to the development system

CHANGE MANAGER provides full management support for the following DB2 data structures:

- storage group
- database
- table space
- table
- check constraint
- foreign key
- indexes
- view
- synonym
- alias

CHECK PLUS

The CHECK PLUS product provides a fast, efficient, and flexible alternative to the IBM DB2 check functions. CHECK PLUS can perform multiple checks in one execution and takes advantage of multiple processors, multitasking, and parallel processing. CHECK PLUS also offers the following functional enhancements:

- provides a comprehensive integrity check of table spaces and indexes to ensure a correct structure and consistency between space maps and data pages
- provides checking of column data through both the VERIFY option and support of table check constraints
- checks the referential integrity (RI) of related tables
- provides for an external definition of referential constraints for applications that cannot afford to define the constraints within DB2
- generates a sequential file that lists foreign keys in RI violation (useful later when correcting the tables)
- provides an SQL batch execution program for deleting rows with foreign key violations
- optionally performs checking while the data is available for read and write purposes with the BMC Software XBM for DB2 SNAPSHOT UPGRADE FEATURE (SUF™)
- provides options for narrowing the range of checking to reduce processing time
- allows for concurrent processing of checking functions
- runs concurrently with other BMC Software utilities on DB2 spaces
- issues diagnostic messages that detail the results of checking processes and allows customizing of a subset of messages
- maintains a table of history statistics about all previous CHECK PLUS activities

CM/PILOT

The CM/PILOT (CM/P) product automates DB2 change management tasks that you normally perform with the CHANGE MANAGER product. With CM/P, you do not need to decide which CHANGE MANAGER processes are required for a task or the sequence in which you need to complete them. Instead, CM/P provides scripts to guide you through the process.

CM/P allows you to choose from 12 predefined scripts or to create your own scripts. You can copy scripts, including those that CM/P provides, and modify them to meet your needs. You can also edit, browse, delete, and perform other maintenance tasks on scripts.

CM/P enables you to utilize the vast change management power of CHANGE MANAGER quickly and easily. By following the dialog panels that are provided in the CM/P scripts or in the scripts that you create, you can specify, analyze, and execute a CHANGE MANAGER worklist to perform the following tasks:

- change data structures
- migrate data structures
- migrate only data
- receive data structure changes
- receive DDL to create data structures
- create full-recovery baselines
- recover data structures with current or old data
- select a user-defined script
- replicate workIDs

When you perform a change management task, CM/P determines the components that CHANGE MANAGER will use and ensures that they run in the correct sequence. CM/P enables you to create tasks that can be done later by someone else or through job scheduling. By reusing the Task ID, you can ensure that the change management task is done the same way every time.

CM/P allows you to easily prepare for various CHANGE MANAGER tasks that process on demand. CM/P dialog panels are task oriented and easy to use. The panels guide you through the CHANGE MANAGER components that are necessary to accomplish specific tasks. The panels prompt you for the information that is needed to build worklists that can process at any time.

CM/P simplifies the decision-making process and enables you to prepare change management tasks that are performed repeatedly. It also lets you create change management tasks that can be performed by less-experienced DBAs. With its unique script-locking feature, CM/P lets you create script steps that cannot be modified and can only be run as you have specified.

CM/PILOT provides the following features:

- automates the change management process
- simplifies change management planning
- allows you to change and migrate multiple objects simultaneously by using an SQL-like data manipulation language (DML)
- lets you create and modify scripts for your environment
- enables you to run change management tasks later, without direct interaction
- enables less experienced DBAs to perform change management tasks
- makes all DBAs more productive

COPY PLUS and C+/MODIFY

The COPY PLUS utility provides a fast, efficient, and flexible alternative to the IBM DB2 COPY utility. Using advanced I/O techniques, COPY PLUS makes copies significantly faster and more efficiently than DB2 COPY. COPY PLUS also offers the following features that reduce manual intervention during the copy process and make the utility easier to use:

- provides the RUNSTATS option to gather and report statistics and to update the DB2 catalog and the BMCSTATS table in the same data pass that is used for the image copy
- supports SUF and XBM, version 3 or later, which provide consistent copies of a group of table spaces while updates are in progress
- produces up to four image copies or DSN1COPY-type copies in a single pass and optionally registers image copies
- provides wildcard selection of table spaces
- provides support for application-owned objects, such as those for SAP R/3 applications, and RECOVERY MANAGER for DB2 groups
- provides grouping and multitasking for copies
- provides facilities for copying DB2 catalog and directory spaces
- allows the duplication of image copies offline
- optionally provides dynamic allocation of copy data sets and performs dynamic tape detection
- provides options for
 - checking the structural integrity of pages
 - establishing a point of consistency before or after the copy process on either a single space or a group of spaces
 - merging a new incremental copy with a previous copy
 - reinstating a previously merged incremental copy
 - automatically escalating incremental requests to full copy requests
 - extending SHRLEVEL control and flexibility
 - compressing disk copies with the BMC Software Extended Compression Architecture® (XCA®) technology
 - controlling the severity of page checking errors
- significantly reduces elapsed time, CPU cycles, and EXCPs for copy processing and optionally optimizes elapsed time for incremental copies
- automatically determines the most efficient I/O access methods and optionally forces a particular access method
- reduces interference with DB2 applications, decreases contention for DB2 resources, and prevents some failures due to utility timeouts

- provides enhanced restart parameter options that reduce manual intervention when restarting a failed copy job
- allows specification of different allocation for the output of full copies that meet a specified size threshold
- is fully compatible with IBM DB2 RECOVER and BMC Software RECOVER PLUS

The C+/MODIFY add-on to the COPY PLUS utility provides a fast, efficient, and flexible alternative to the IBM DB2 MODIFY utility. C+/MODIFY is a SYSIBM.SYSCOPY table management utility that provides the DBA with a wide range of functionality for the maintenance of the information in that table. Maintenance of the COPY PLUS BMCXCOPY table is also possible using C+/MODIFY.

C+/MODIFY provides the following functionality to save time and resources:

- provides a more granular method and more options for deleting SYSCOPY rows
- provides maintenance of the BMC Software BMCXCOPY table, which handles copy registration for index spaces that COPY PLUS has copied
- supports wildcard selection, giving DBAs a quick mechanism for setting up jobs and eliminating the need to maintain the jobs when spaces are created or dropped
- provides support for RECOVERY MANAGER for DB2 groups and SAP R/3 application-owned objects
- allows registration of image copies not produced by the COPY utility, which is useful in some drop recovery and disaster recovery backup scenarios
- supports modification of the site information by allowing you to alter the ICBACKUP information that designates the site type, the order of use of the copy, and the DEVTYPE information for your recovery site
- improves performance by enhancing the granularity of deletions and providing the ability to commit more often, which alleviates problems than can occur if you neglect regular maintenance of SYSCOPY
- synchronizes the ICF catalog and SYSCOPY by analyzing and removing rows for which there is no actual data set and by providing a utility that optionally cleans up the Intersystem Communications Facility (ICF) catalog at the same time
- verifies that table spaces are recoverable by analyzing SYSCOPY and by allowing you to designate some rules for recoverability based on the number of copies, the number of logs, or the number of days
- works with COPY PLUS to make copies when objects are found to be not recoverable

Cross-System Image Manager

The Cross-System Image Manager (XIM) technology provides sysplex performance improvements by enabling the distribution and management of discrete units of work (UOW) across one or more MVS systems. BMC Software utility products that exploit XIM can divide single, long-running tasks into multiple parallel tasks to be run across multiple machines in the sysplex, thus decreasing the overall elapsed time.

XIM provides the following key features:

- automatically distributes across multiple MVS images the workload that was previously performed in a single utility step
- performs dynamic and reactive workload balancing based on processor speed and resource usage

XIM provides the following benefits:

- reduces the amount of time required to complete utility functions
- increases application availability by minimizing the time your applications are offline
- optimizes the use of hardware, software, and personnel through workload balancing
- lets you exploit parallel processing for database utility tasks

XIM delivers parallel enablement for the BMC Software UNLOAD PLUS for DB2 utility product.

DASD MANAGER PLUS

The DASD MANAGER PLUS product is a comprehensive database management tool that supports DB2 version 3.1 and later. DASD MANAGER PLUS

- gathers, stores, and analyzes statistics
- monitors changes in the database
- automates utility execution based on the condition of the data
- estimates space requirements
- produces numerous analytical and statistical reports about DB2 objects
- allows you to specify and generate IBM and BMC Software utility jobs, including the required JCL

DASD MANAGER PLUS gathers and analyzes statistics with BMCSTATS. BMCSTATS collects all of the statistics that RUNSTATS and STOSPACE collects, plus many other statistics that can be helpful in administering a database. BMCSTATS maintains these statistics in a historical database. The historical database makes possible trend analysis of DB2 object activity. If you use LOADPLUS, REORG PLUS, or COPY PLUS, these utilities optionally provide statistics to the BMCSTATS database.

DASD MANAGER PLUS automates utility operation through an analysis tool, BMCTRIG. BMCTRIG examines the BMCSTATS database for exception conditions, such as number of extents. When conditions reach or exceed the trigger values that you specified, DASD MANAGER PLUS can automatically build and submit the appropriate utility job to correct the condition. BMCTRIG has built-in intelligence that eliminates redundant work and matches object types with utilities.

The Space Estimation function simulates changes to the database. You can use this function to predict how changes will affect page size, segment size, primary and secondary allocations, number of rows, index type, average row length, and so on.

DASD MANAGER PLUS provides a full-function interface for entry, validation, and generation of DB2 utilities and commands. Additionally, you can save and reuse each user-specified set of utility parameters.

This product provides numerous statistical and analytical reports. Many of the reports are written in REXX and are customizable. You can run the customizable reports as they are written, or edit and run them without compiling, linking, and binding the code. You can also create your own REXX reports by using the external functions provided.

DATA PACKER

The DATA PACKER product reduces Direct Access Storage Device (DASD) requirements for DB2 tables by compressing table rows. DATA PACKER compresses and expands table rows through the use of the DB2 EDITPROC exit facility. DATA PACKER makes no modifications to DB2 or application programs.

The BMC Software exclusive XCA technology provides common compression routines used by DATA PACKER. The architecture ensures data integrity during table compression and expansion.

The features of DATA PACKER allow you to

- reduce DASD requirements for DB2 tables by as much as 80 percent
- decrease the number of physical I/Os needed (more rows will fit in a DB2 table space page)
- perform trial compression on uncompressed tables
- choose from multiple compression techniques (including hardware-assisted)
- improve database response time and performance
- reduce requirements for DB2 database buffer space
- make faster image copies that require less storage media
- improve DASD capacity planning by supplying compression statistics
- add an EDITPROC to a table without unloading, dropping, and reloading the table
- dynamically change EDITPROCs without unloading, dropping, and reloading the table
- add columns to a table that has an EDITPROC defined, without unloading, dropping, and reloading the table

EXTENDED BUFFER MANAGER

The EXTENDED BUFFER MANAGER (XBM) product provides system-wide performance improvement and increases DB2 data availability. Also available are XBM products that provide performance improvement for CICS, IMS, and MQ, and data availability for CICS, IMS, and VSAM.

- *Performance:* XBM improves system-wide performance by reducing the number of physical read I/Os performed in accessing data. XBM satisfies read requests by caching frequently accessed data using central or expanded storage.
- *Availability:* The XBM Snapshot Utilities component (available separately as the SNAPSHOT UPGRADE FEATURE), when used with the BMC Software COPY PLUS, CHECK PLUS, REORG PLUS, or UNLOAD PLUS product, allows you to run these utilities while DB2 data remains available to users. The Snapshot Utilities component uses the XBM caching technology to store point-in-time copies. This feature creates a point of consistency and releases the data, allowing updates to occur while utility operation continues. Also, you can use the XBM Storage Systems Integration (SSI) component with the BMC Software COPY PLUS product and supported intelligent storage devices to provide point-in-time copies without requiring a software cache.

XBM provides the following benefits:

- supports both I/O caching and Snapshot Utilities processing in DB2 data sharing environments
- provides completely dynamic allocation and manipulation of the extended buffer to improve immediate system I/O requirements
- compresses data stored in cache to optimize system resources
- improves read access times for sequential I/O operations
- optimizes database performance of DB2 objects at the data set level
- features a menu-driven ISPF interface to facilitate setting up, activating, and monitoring all DB2 I/O activity
- provides an online monitor for DB2 objects, Snapshot Utilities activity, and individual cache utilization statistics
- uses any combination of fixed virtual storage, pageable virtual storage, data space, or ESO hiperspace for cache
- accommodates both 4K and 32K DB2 page types in the same cache

OPERTUNE

The OPERTUNE product allows you to modify DB2 subsystems to resolve performance tuning or operations management problems. Using the simple OPERTUNE command syntax, you can change a DB2 ZPARM value dynamically. You issue the commands by using the OPERTUNE ISPF interface, from a batch job step, or from the console. You can also define a set of ZPARM values to support a specific processing environment and then schedule the set to be issued to DB2 automatically. The schedule can be issued based on time of day and day of the week, or on demand.

OPERTUNE performs these changes without bringing the DB2 subsystem down or logging users off of the system. A RESET command can restore individual DB2 parameters to their original value or return all DB2 parameters modified by OPERTUNE to their original values.

You can secure OPERTUNE by using its internal security or a standard security package (such as RACF or ACF2). Security can be controlled at the user, subsystem, or feature (individual ZPARM) level.

With OPERTUNE, you can

- modify DB2 parameters to relieve current problems and bottlenecks
- dynamically tune DB2 subsystems
- adapt resource assignment and allocation to workload fluctuations
- enhance 24 hour by 7 day continuous operations for DB2
- enhance automated operations for DB2
- dynamically add and remove active logs
- cancel a thread that is draining resources without cancelling the thread's address space
- free a table space or partition of all users
- tune remote DB2 subsystems through VTAM or by using XCF groups
- test applications with different DB2 parameters without recycling DB2
- accept commands from the BMC Software ACTIVITY MONITOR product when exceptions are detected, from the ACTIVITY MONITOR Command Interface panel, or from ACTIVITY MONITOR reports
- submit commands to the BMC Software XBM utility from the Free Form Commands panel, or combine XBM commands into groups and submit them automatically through OPERTUNE schedules

PACLOG

The PACLOG product provides substantial archive log media savings along with a set of log management functions that allows you to optimize the number of active logs and their sizes. PACLOG provides the following features:

- substantial resource savings by reducing the amount of media required for archive logs by up to 95 percent
- optional data compression, whether writing to tape or disk

PACLOG uses proven BMC Software XCA technology when writing to disk.

- removal of log data not needed for any DB2 recovery process

PACLOG optionally removes index-related log records, record types not needed by any DB2 forward recovery process, and log data for user-specified DB2 objects.

- a modeling tool for optimization of active log size
- multiple copies of archive logs

PACLOG can create up to four copies of the original log: two for local site use and two for recovery site use.

- media, filtering, and compression options that may be separately specified for each copy
- stacking of copies to tape
- simulation of the archive log process, with no impact on production
- log management reports that notify you when an archive log is no longer needed and when an archive log has been processed
- dynamic allocation of input and output archive log data sets

RECOVER PLUS

The RECOVER PLUS product is a reliable, functionally enhanced alternative to the IBM DB2 RECOVER utility. By providing additional recovery strategies, maximized concurrency, and additional options, RECOVER PLUS can recover table spaces or index spaces significantly faster and at less cost than DB2 RECOVER. RECOVER PLUS has the following features:

- recovers a space to the current state or to a prior point
- recovers index spaces from image copies and logs, or rebuilds them from data in the table space
- allows sorting of log records, optionally reading several log files concurrently
- provides a choice of recovery strategies, including BACKOUT
- supports the data set level recovery of index spaces
- allows keywords LASTCOPY, LASTQUIESCE, LASTCOMMONQ, LASTARCHQ, and LASTSHUTDOWN for partial recoveries
- allows image or DSN1COPY-type copies to be made during processing (when logs are sorted)
- dynamically allocates input image copies, output image copies (optional), and log data sets
- allows extraction of index keys during table space recovery and allows them to be sent directly to a sort task as they are extracted
- allows keys for nonpartitioning indexes to be unloaded in concurrent runs. The keys may be sorted during the UNLOADKEYS runs or during a later run, which builds the complete index.
- allows concurrent recovery of the partitions of a partitioned table space, or of the data sets of a nonpartitioned space
- provides restart capabilities
- allows preview of recovery activity and resource usage
- detects stacked tape input, orders activities optimally, and keeps tapes mounted until all the needed data sets are read
- allows you to control resource usage by limiting the number of logs to be read simultaneously and by limiting the number of tape drives to be used
- provides a list of archive tapes and copy tapes that are required for recovery
- provides early recall of migrated data sets
- supports fallback processing
- allows optional reallocation of STOGROUP-defined table space and index space data sets
- allows the redirection of recovery output to non-DB2 data sets
- supports DB2 data sharing

R+/CHANGE ACCUM

The R+/CHANGE ACCUM product offers the ability to create new recovery resources. Using this capability in conjunction with RECOVER PLUS can significantly streamline normal recovery and disaster recovery processes.

R+/CHANGE ACCUM extracts and sorts updates from the DB2 log for a designated group of objects and stores the updates in a file. Functions for R+/CHANGE ACCUM use the same JCL as RECOVER PLUS and are fully integrated with RECOVER PLUS.

R+/CHANGE ACCUM is composed of an interactive ISPF interface and a batch utility. You can use the interface to define and create change accumulation groups. A change accumulation group can include any number of table spaces, partitions, or data sets, as long as the EXEC statement specifies a sufficient amount of REGION. R+/CHANGE ACCUM stores group definitions in the R+/CHANGE ACCUM repository. You can use the R+/CHANGE ACCUM batch utility to generate change accumulation files routinely.

R+/CHANGE ACCUM has the following features:

- provides change accumulation files, which are a recovery resource not available in DB2, while increasing availability of DB2 data
- extracts and sorts DB2 log records prior to a recovery to reduce the time spent processing DB2 log records during a recovery that uses RECOVER PLUS
- supports the extraction of log records for indexes
- supports the concurrent reading of several log files
- reduces the elapsed time required to recover table spaces and indexes by providing greater concurrency
- extracts only log data needed for recovery, and only for selected spaces
- allows selected log data to be kept online and compressed with any compression product
- performs change accumulation and recovery within the same job step, with only one pass of the DB2 log
- provides a method for preparing for recovery without accessing the spaces
- provides a way to group table spaces to allow change accumulation of multiple table spaces on the same file
- allows the accumulation of updates for multiple groups of DB2 objects in one pass of the log
- supports DB2 data sharing
- reduces the number of tape mounts required for a recovery in a data sharing environment

R+/RESOURCE MAXIMIZER

The R+/RESOURCE MAXIMIZER product allows you to utilize recovery resources in new ways. R+/RESOURCE MAXIMIZER syntax works with base RECOVER PLUS statements to expand the use of the DB2 log, copies, and R+/CHANGE ACCUM change accumulation files, thereby enabling new capabilities.

Functions for R+/RESOURCE MAXIMIZER use the same JCL as RECOVER PLUS and are fully integrated with RECOVER PLUS. With R+/RESOURCE MAXIMIZER, you can

- recover a dropped table space
- use OUTCOPY, INCOPY, and OBID translation features to support migrating data between DB2 table spaces on the same system or between different sites
- create copies from prior copies and the DB2 log
- use copies that are not registered in SYSIBM.SYSCOPY
- translate internal DB2 object IDs in the DB2 table space as it is processed
- read the log without referencing the DB2 SYSIBM.SYSLGRNX table during recovery
- run in a DB2 data sharing environment
- use the R+/RESOURCE MAXIMIZER copy analyzer, ADFCOPY, to evaluate the data on an image copy data set

RECOVERY MANAGER for DB2

The RECOVERY MANAGER for DB2 (RMGR) product allows you to create groups of DB2 table spaces and indexes and then automatically generate JCL to back up or recover those groups using BMC Software utilities or IBM utilities. You can also maintain and recover a wide range of system resources. RECOVERY MANAGER provides the following features and functions:

- allows groups to be built interactively starting at the table space, index, volume, storage group, plan, or package level; allows volume groups to be built in batch mode; allows automatic creation of groups for an entire DB2 subsystem (especially useful for large applications such as SAP R/3)
- supports transaction recovery in conjunction with the BMC Software Log Master and Apply Plus products
- supports the backup and recovery of an entire DB2 subsystem
- supports backup and recovery avoidance of unchanged objects, thus increasing efficiency of the backup and recovery process
- supports two different levels of hardware mirroring
- allows wildcard generation of object lists
- allows groups to be combined to form new groups
- allows volume recovery processing to be performed after a failure
- supports standard DB2 disaster recovery
- supports a wide range of COORDINATED RECOVERY MANAGER (CRM™) for MVS functions
- maintains consistency of related objects; checks RI and indexes and supports recovery, repair, and check actions
- generates accurate, optimized, multiple job recovery JCL online or in batch mode; sizes data sets and generates data definition (DD) cards for stacked tape
- provides a repository for storing recovery plans
- allows recovery type selection (to current, copy, quiesce, or relative byte address (RBA))

- provides a range of recovery and maintenance of system resources (catalog and directory, active logs, archive logs, boot strap data set (BSDS)) for local sites and for disaster recovery
- allows recover utility syntax option selection at the object level, group level, and subsystem level
- provides range of group user authorizations, including PUBLIC
- provides a choice of recovery strategies: supports LOGONLY recoveries and the RECOVER PLUS LOGSORT strategy
- optionally makes image copies after recovery; supports tape stacking, symbolic data set names, and generation data groups (GDGs)
- provides a choice of utilities—RECOVER PLUS (including sorted UNLOADKEYS support) or DB2 RECOVER, COPY PLUS or DB2 COPY, CHECK PLUS or DB2 CHECK
- uses a menu-driven user interface that includes online help for task panels and message help

RECOVERY MANAGER for OS/390

The RECOVERY MANAGER for OS/390 product automates the backup and recovery of OS/390 data sets in both sysplex and nonsysplex environments and ensures fast and successful task completion. RECOVERY MANAGER lets you plan for various recovery scenarios, including volume failure and disaster recovery. RECOVERY MANAGER features allow you to

- identify the data sets that an application uses
- predefine application groups (group all data sets belonging to an application)
- define appropriate backup and recovery options for each application group
- generate full volume backups and restores
- revalidate the definition of your groups
- recover tape catalogs to a backup or a point in time
- recover ICF catalogs to a backup or a point in time
- audit the recoverability of all objects
- report on availability of data sets needed for recovery
- resynchronization of the recovery assets catalog (RASCAT) and the repository
- build a DFDSS stand-alone IPL tape
- build an ICKDSF stand-alone IPL tape
- enforce recovery rules
- create and save recovery plans to be used for disaster recovery
- test your recovery plans before failures occur

These features combine to provide solutions to most recovery management problems, including those that involve cost, speed, and accuracy. RECOVERY MANAGER provides the following major benefits:

- reduces costs and saves time in initiating a recovery because of the immediate availability and readiness of a plan that is already tailored to fit the recovery scenario
- reduces costs and time savings in all backup and recovery operations due to the product's ease of use and its ability to automatically generate accurate optimized JCL that guarantees success
- reduces costs and saves time in disaster recovery by providing
 - ability to create and save disaster recovery plans
 - support for contingency planning at the local site
 - offsite copies of application data sets
 - JCL for recovering applications at a recovery site
 - verification of resources required at the recovery site

RECOVERY MANAGER provides the additional, equally important but intangible benefit of peace of mind for the system programmer.

REORG PLUS

The REORG PLUS utility provides a fast, efficient, and flexible alternative to the IBM DB2 REORG utility. It does so, in part, by taking advantage of multiple processors, multitasking, and parallel processing. REORG PLUS accomplishes standard reorganization tasks and also offers the following features:

- conditionally performs a reorganization based on DB2 statistics for the object
- works with the UTILITY MANAGER for DB2 product to automate the reorganization process
- analyzes resources for the specific reorganization, provides data about resource requirements, and optionally dynamically allocates your work files (based on the analysis)
- reorganizes any number of selected partitions of a table space and index space
- offers an optional single processing phase for better performance
- optionally purges unwanted or out-of-date information, and allows you to write this data to an archive data set
- optionally updates column values to a specified constant value
- restores PCTFREE and FREEPAGE space in nonleaf pages of indexes
- builds or keeps a compression dictionary when the object being reorganized has the COMPRESS YES attribute
- optionally collects statistics for the BMCSTATS database (if the BMC Software DASD MANAGER PLUS product is installed) and the DB2 catalog
- produces up to four standard or inline copies or DSN1COPYs concurrently with the reorganization and optionally dynamically allocates the data sets
- optionally deletes and redefines user-defined data sets and data sets defined in DB2 storage groups as part of the reorganization
- optionally allows the object being reorganized to be available in read-only (RO) status during unload and reload processing

- maintains a table of history statistics on all previous REORG PLUS activities
- supports the AMEND option of the BMC Software DATA PACKER product
- runs concurrently with other BMC Software utilities on DB2 spaces
- supports double-byte character set (DBCS) characters

You can also perform an online reorganization by specifying SHRLEVEL CHANGE on the command. The online reorganization allows read and write application access to DB2 objects during most of the reorganization process and significantly decreases the amount of time the data is unavailable.

REORG PLUS with SHRLEVEL CHANGE offers the following features:

- supports DB2 versions 6, 5, 4, and 3
- runs outside of DB2 so the log apply processing does not impact DB2 buffer pools and other resources
- does not cause -904 SQLCODES
- creates a row ID (RID) translation map that is not a DB2 object and does not use DB2 resources
- optionally purges unwanted or out-of-date information and allows you to write this data to an archive data set during an online reorganization
- provides multitasking of the process of switching the staging data sets to the original data set names
- does not require the mapping table to be predefined, and does not require the mapping table name to be specified on the REORG command
- dynamically allocates the image copy data sets, the staging data sets, and the spill data sets, and allows you to specify the size of the spill data sets
- allows you to alter processing dynamically by using the Utility Monitor function of XBM or the SUF component of XBM
- provides an interface for automated operations by allowing you to use the XBM or SUF MVS console support to alter operation

RESOLVE Log Master for DB2

The RESOLVE Log Master for DB2 product is an important, useful tool that unlocks the resources embedded in your DB2 logs. With this tool, you can easily access your logs and use their information for purposes beyond ordinary restart or recovery. You can access logs while DB2 is up and running, without diminishing database performance. RESOLVE Log Master can make your DB2 environment more productive in the following ways:

- enables logical backout of problem transactions, by generating UNDO SQL or REDO SQL statements for specific transactions
- analyzes problem transactions, with unique backout integrity checking to identify when backout processing will affect updates performed later
- supports data migration with MIGRATE SQL or MIGRATE DDL statements and other features
- enables database auditing without traces or coding of new applications
- provides automated recovery of DB2 objects that have been dropped from the DB2 catalog (using UNDO DDL and working with the BMC Software product R+/RESOURCE MAXIMIZER)
- processes log records for DB2 objects that are no longer defined in the DB2 catalog
- completes partial log records from either the table space or image copies
- supports DB2 security at the table level, with options to suppress log records in output when a job lacks required authorization
- works with storage management software to recall migrated data sets and migrate them back to their original level
- provides extensive filtering capabilities to easily select only the necessary log records
- provides a comprehensive set of reports, with flexible sorting criteria and the ability to include or exclude columns
- provides the convenient logical log output format, designed for application programs with easy-to-use fields in a published format
- includes an integrated SQL processor that supports customized responses to SQL return codes and the use of alternate indexes
- analyzes the DB2 log to locate quiet points for specific tables or table spaces, enabling more precise recovery

UNLOAD PLUS

The UNLOAD PLUS product is a high-performance utility for unloading DB2 data. UNLOAD PLUS also offers the following features:

- unloads rows from one or more DB2 tables (including DB2 catalog tables) in one or more table spaces with one utility invocation
- provides powerful selection criteria for unloading specific rows
- uses I/O, CP and sysplex parallelism
- allows you to tailor multitasking activities by starting more than one task per CPU for individual or all processing phases of UNLOAD PLUS, thereby improving performance
- dynamically eliminates processing of partitions that do not meet the conditions of the WHERE clause
- optimizes the data retrieval of single tables in a multi-table segmented table space by processing only those segments that pertain to the table that is being unloaded. UNLOAD PLUS does not read data in tables that were not specified by the selection criteria.
- provides for unloading from full or incremental image copies, DSN1COPY data sets, INLINE copies, or VSAM linear data sets
- optionally produces multiple unload data sets either by multiple SELECT statements or by partition using a single SELECT statement
- provides a comprehensive set of data type conversions and a user exit facility for any conversions not provided
- optionally allows UNLOAD PLUS to use DDL for DB2 object definitions, instead of using object definitions in the DB2 catalog, when unloading using the INFILE *ddname* option
- optionally sorts the output records by specified columns or clustering key; for multiple table unloads, sorted output records are also grouped by table
- provides an option for unloading a sample of rows from a table or table space
- provides an option for limiting the number of rows unloaded from a table or table space

- optionally provides utility-generated or user-specified output-record tagging with a character string constant
- allows replacement of NULL values with a user-specified constant or automatically generates a null indicator value with each output field for each nullable column selected
- optionally produces the LOAD control statements needed for loading the data back into a DB2 table, or file definition statements for use with FOCUS, SAS, EASYTRIEVE PLUS, Teradata, SQL/DS, or NOMAD

UTILITY MANAGER

The UTILITY MANAGER product is a robust “out-of-the-box” automation solution that allows you to create the optimal solution in your environment. Many of the components that constitute UTILITY MANAGER are provided in source in REXX. You may customize these components to fit your requirements. UTILITY MANAGER also offers the following features and benefits:

- simplifies database maintenance as objects become larger and more numerous

UTILITY MANAGER works with other BMC Software products to support an automated utility batch cycle that performs required maintenance.

- reduces the daily burden of tailoring JCL and utility commands to maintain the growing number of database objects
- conserves system resources and allows you to profit from the flexible maximum use of shrinking batch windows

UTILITY MANAGER allows you to define processing windows to indicate when the utilities are allowed to operate and which objects are affected. You can submit multiple batch jobs to run in parallel, increasing the throughput of the batch window.

- provides three different modes of performing object registration for utility automation: manually, using a batch job, or automatically

Manual object registration is performed through the BMC Software CATALOG MANAGER for DB2 product or by the BMC Software Catalog Browser interface. Batch registration is accomplished by running the BMC Software catalog scan (UTA\$SCAN) sample program supplied with UTILITY MANAGER. Automatic registration is performed by an interface to MV/Alarm and AutoOPERATOR.

- allows you to modify the BMC Software catalog scan (UTA\$SCAN) sample program to meet the needs of your installation.

This product evaluates DB2 catalog statistics and registers objects whose statistics meet designated criteria.

Index

A

- ACF2 32
- ACTIVITY MONITOR 38
- ALTER 39
 - ALTER authority 33
 - ALTER INDEX authority 33
- ANALYZE processing phase 25
- Apply Plus 41
- APPTUNE 42
- ASCII character data 3
- authority
 - ALTER 33
 - ALTER INDEX 33
 - CONTROL 33
 - DB2 32
 - DBADM 32
 - DBCTRL 32
 - LOAD 32
 - SYSADM 32
 - SYSCTRL 32
 - UPDATE 33
- authorization mechanisms 32
- authorizations needed to execute LOADPLUS
 - DB2 32
 - system (RACF) 33

B

- benchmark tests 7 to 10
- benefits of LOADPLUS 2, 5 to 6

- BMC Software products for DB2 and OS/390
 - ACTIVITY MONITOR 38
 - ALTER 39
 - Apply Plus 41
 - APPTUNE 42
 - C+/MODIFY 48
 - CATALOG MANAGER 43
 - CHANGE MANAGER 44
 - CHECK PLUS 45
 - CM/PILOT 46
 - COPY PLUS 48
 - Cross-System Image Manager (XIM) 50
 - DASD MANAGER PLUS 51
 - DATA PACKER 52
 - EXTENDED BUFFER MANAGER (XBM) 53
 - OPERTUNE 54
 - PACLOG 55
 - R+/CHANGE ACCUM 57
 - R+/RESOURCE MAXIMIZER 58
 - RECOVER PLUS 56
 - RECOVERY MANAGER for DB2 59
 - RECOVERY MANAGER for OS/390 61
 - REORG PLUS 63
 - RESOLVE Log Master for DB2 65
 - RMGR 59
 - UNLOAD PLUS 66
 - utilities list 37
 - UTILITY MANAGER 68
 - XBM 53
 - XIM 50
- BMCCPYnn data sets 30
- BMCCPZnn data sets 30
- BMCDICT table 23, 26

BMCHIST table 23, 26
BMCRCYnn data sets 31
BMCRCZnn data sets 31
BMCSYNC table 23, 26
BMCUTIL table 23, 26
buffering 24

C

C+/MODIFY 48
CA-ACF2/DB2 32
CATALOG MANAGER 43
CA-Top Secret/DB2 32
CHANGE ACCUM 57
CHANGE MANAGER 44
character data
 ASCII 3
 EBCDIC 3
CHECK PLUS 45
CM/PILOT 46
COMBINED processing phase
 considerations 34
 primary functions 25
command syntax, examples 11 to 16
comparison of LOADPLUS to IBM's DB2
 LOAD utility 3
compression (hardware), how LOADPLUS uses
 and performs compression 36
CONTROL authority 33
conventions
 document xvii
 typographical xvii
conversions, allowable data types 19
copy data sets 30
COPY PLUS 48
Cross-System Image Manager (XIM) 50

D

DASD MANAGER PLUS 51
DATA PACKER 52
data sets of LOADPLUS
 BMCCPYnn 30
 BMCCPZnn 30
 BMCRCYnn 31
 BMCRCZnn 31

copy data sets 30
DB2 31
multiple 24
SORTOUT 24, 29
SORTWKnn 29
SYSDISC 30
SYSERR 30
SYSIDCIN 30
SYSIN 29
SYSPRINT 31
SYSREC 24, 29
SYSUT1nn 29
UTPRINT 31

data type conversions table 19

data, character 3

DB2

 authorizations required to execute

 LOADPLUS 32

 data sets 31

 LOAD utility compared to LOADPLUS 3

 LOADPLUS interaction with 23

 use of resources 23

 versions required for running LOADPLUS
 22

DBADM authority 32

DBCTRL authority 32

DCI install system 22

document conventions xvii

documentation

 online xv

 related xv

 release notes xvi

dynamic allocation

 benefits 6

 data sets affected 29 to 31

E

EBCDIC character data 3

environment for operating LOADPLUS 22

examples, LOADPLUS processing 11 to 16

EXTENDED BUFFER MANAGER (XBM) 53

F

features of LOADPLUS 5 to 6

H

hardware compression, how LOADPLUS uses
and performs compression 36

I

I/O optimization 24
image copies 3
inline image copies 3
Install System (DCI) 22
installation 22
interaction with DB2 23

L

listing of BMC Software utilities 37
LOAD authority 32
LOAD processing phase 25
LOAD utility, of DB2, compared to
LOADPLUS 3

M

maximum performance, how LOAPLUS
achieves 22
messages, online MVS data set xv
multiple data sets 24
multiprocessing environment 23
multitasking 3, 24

O

online documentation xv
operating environment 22
OPERTUNE 54
optimization of I/O 24

P

PACLOG 55
performance
 benchmarks 7
 maximum, how LOADPLUS achieves 22
 tuning options 24
phases of LOADPLUS
 ANALYZE 25
 COMBINED 25
 LOAD 25
 PRELOAD 25
 single-phase load diagram 28
 two-phase load diagram 27
 UTILINIT 25
 UTILTERM 25
PRELOAD processing phase 25
processing examples 11 to 16
processing phases
 ANALYZE 25
 COMBINED 25
 LOAD 25
 PRELOAD 25
 single-phase load diagram 28
 two-phase load diagram 27
 UTILINIT 25
 UTILTERM 25

R

R+/CHANGE ACCUM 57
R+/RESOURCE MAXIMIZER 58
RACF
 authorizations required to execute
 LOADPLUS 33
 security exit 32
RECOVER PLUS 56
recoverability of loaded table space 36
RECOVERY MANAGER for DB2 (RMGR) 59
RECOVERY MANAGER for OS/390 61

referential integrity issues 35
related publications xv
release notes xvi
REORG PLUS 63
RESOLVE Log Master for DB2 65
RESOURCE MAXIMIZER 58
resources, how LOADPLUS saves them 7
RMGR 59

S

secondary authorization ID 32
security 32
single-phase load
 considerations 34
 diagram of phases 28
 primary functions 22
SORTOUT data sets 24, 29
SORTWKnn data sets 29
SQL, use of 23
storage, virtual 23
syntax, examples 11 to 16
SYSADM authority 32
SYSCTRL authority 32
SYSDISC data sets 30
SYSERR data sets 30
SYSIDCIN data sets 30
SYSIN data sets 29
SYSPRINT data sets 31
SYSREC data sets 24, 29
system authorizations needed to execute
 LOADPLUS 33
SYSUT1nn data sets 29

T

table space recoverability 36
tables used by BMC Software utilities
 BMCDICT 23, 26
 BMCHIST 23, 26
 BMCSYNC 23, 26
 BMCUTIL 23, 26
time, how LOADPLUS saves it 7
TopSecret 32
tuning options 24

two-phase load
 diagram of phases 27
 primary functions 22
typographical conventions xvii

U

UNLOAD PLUS 66
UPDATE authority 33
UTILINIT processing phase 25
utilities list 37
UTILITY MANAGER 68
UTILTERM processing phase 25
UTPRINT data sets 31

V

virtual storage 23


W

work files, dynamic allocation
 benefits 6
 data sets affected 29 to 31

X

XBM 53
XIM 50

Notes



100031493